

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : G06F	A2	(11) International Publication Number: WO 00/36491 (43) International Publication Date: 22 June 2000 (22.06.00)
<p>(21) International Application Number: PCT/US99/30274</p> <p>(22) International Filing Date: 15 December 1999 (15.12.99)</p> <p>(30) Priority Data: 60/112,817 17 December 1998 (17.12.98) US</p> <p>(71) Applicant (for all designated States except US): CALIFORNIA INSTITUTE OF TECHNOLOGY [US/US]; 1200 East California Boulevard, Mail Code 210-85, Pasadena, CA 91125 (US).</p> <p>(72) Inventors; and (75) Inventors/Applicants (for US only): THORNLEY, John [US/US]; 1200 East California Boulevard, Mail Code 256-80, Pasadena, CA 91125 (US). CHANDY, K., Mani [US/US]; 1200 East California Boulevard, Mail Code 256-80, Pasadena, CA 91125 (US). ISHII, Hiroshi [US/US]; 1200 East California Boulevard, Mail Code 201-85, Pasadena, CA 91125 (US).</p> <p>(74) Agent: HARRIS, Scott, C.; Fish & Richardson, P.C., Suite 1400, 4225 Executive Square, La Jolla, CA 92037 (US).</p>		<p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published Without international search report and to be republished upon receipt of that report.</p>

(54) Title: PROGRAMMING SYSTEM AND THREAD SYNCHRONIZATION MECHANISMS FOR THE DEVELOPMENT OF SELECTIVELY SEQUENTIAL AND MULTITHREADED COMPUTER PROGRAMS

(57) Abstract

A structured multithreaded programming system is described for integrated use with existing and new programming languages and systems. The structured multithreaded programming system enables programs to be developed which include both multithreaded and multithreadable code constructs. The multithreaded code constructs require explicitly concurrent execution. The multithreadable code constructs can be executed either sequentially or concurrently, at the selection of the programmer or computer user. When executed concurrently, the different threads of execution in a multithreaded program developed with this system can be synchronized using innovative synchronization objects. One type of synchronization object is a special type of counter, which can be constrained to be monotonically increasing in value. Another related type of synchronization object is a special type of flag, which can be constrained to have its value set monotonically.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakhstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

WO 00/36491

PCT/US99/30274

PROGRAMMING SYSTEM AND THREAD SYNCHRONIZATION MECHANISMS
FOR THE DEVELOPMENT OF SELECTIVELY
SEQUENTIAL AND MULTITHREADED COMPUTER PROGRAMS

5

The present application claims priority under 35
U.S.C. 119(e) from provisional application number
60/112,817 filed December 17, 1998.

10

Background

Many computer programs are computationally
intensive, meaning that they require large amounts of
15 computing power. As a consequence, these programs may
execute more slowly than the computer user desires, even
on the fastest computers. One way of increasing the
execution speed of a computationally intensive computer
program is to divide the program into multiple units, or
20 loci, of concurrent execution. These units of execution
are known as "threads". A program with multiple threads
of execution is known as a "multithreaded program". A
program with only a single thread of execution is known
as a "sequential program". The threads that make up a
25 multithreaded program may be executed concurrently on
multiple computer processors, allowing many operations in
the program to be carried out simultaneously, thereby
speeding up program execution.

WO 00/36491

PCT/US99/30274

In a multithreaded program, the program or operating system must control the access of threads to data objects in the program, in order to prevent the multiple threads from concurrently accessing the same data object in an undesirable manner. If multiple threads modify the same data object concurrently, or read and modify the same data object concurrently, the resulting state of the program is extremely difficult to determine. Developing a multithreaded program is significantly more difficult than developing a sequential program because of the problems of (1) expressing the division of a program into multiple threads and (2) structuring and controlling the access of those threads to data objects.

15

Summary

The present application teaches a structured thread ("Sthread") system with thread synchronization and production mechanisms.

Another aspect produces multithreadable code. The multithreadable code can be annotated using information indicative of its multithreadability. The multithreadable code constructs are code constructs that can be executed in a multithreadable manner, or

WO 00/36491

PCT/US99/30274

equivalently in a sequential manner. Multithreadable code constructs may be expressed by annotating sequential code constructs to indicate that their multithreaded execution is equivalent to sequential execution. The

5 multithreadable code can be used along with multithreaded code. Specific instances of a multithreadable constructs: a multithreadable block, and a multithreadable for loop, and are disclosed.

The second aspect of the system is the integration

10 of multithreadable code constructs with traditional explicitly multithreaded code constructs. Explicitly multithreaded code constructs must always be executed in a multithreaded manner. Examples of explicitly multithreaded code constructs include multithreaded block

15 constructs, multithreaded for loop constructs, and library-based thread creation functions. Multithreadable code constructs and explicitly multithreaded code constructs may be intermingled within a program as required, with well-defined meaning.

20 According to a first aspect, a special counter called an "s-counter", is used as a thread synchronization mechanism. Special "s-Flags" can also be used for thread synchronization, and flag synchronization is also described herein.

WO 00/36491

PCT/US99/30274

Yet another aspect is the implementation of the programming system within an existing compiler environment using a special pre-processing system.

The embodiments of the invention describe additional
5 details, including the following:

The s-counter synchronizing the access of threads to shared data objects. The mechanisms use "monotonic" synchronization objects, with operations that can be constrained to only move the value of the object in one
10 direction. Monotonic synchronization objects can be used to synchronize the access of threads to shared data objects in multithreadable code constructs in a manner that guarantees the equivalence of sequential and multithreaded execution. Specific instances of monotonic
15 synchronization objects are disclosed, namely a form of counter called an "s-counter" and a form of flag called an "s-flag". The s-counter is a particularly powerful thread synchronization mechanism in many contexts, with its use in multithreadable code constructs being one
20 example.

The application describes implementation of the multithreaded programming system within an existing program development and compilation environment using a special source-to-source preprocessing system and high-

WO 00/36491

PCT/US99/30274

level thread library. This allows the system to be transparently and seamlessly integrated with existing programming systems such as Microsoft Visual Studio for the Microsoft Windows family of operating systems, the
5 GNU program development tools for Linux and other versions of the Unix operating system, or on any version of the Java programming language, for example.

Brief Description Of The Drawings

10 These and other aspects will now be described in detail with respect to the accompanying drawings, wherein:

Figure 1 is a process flowchart showing a prior art method for compiling multithreaded code;

15 Figure 2 shows a computer system and its thread allocation system;

Figure 3 shows a flowchart of defining an s-counter; and

Figure 4 shows a flowchart of Sthreads execution of
20 a program.

Detailed Description

FIG. 1 is a process flowchart showing a prior art method for compiling multithreaded code. Source code text

WO 00/36491

PCT/US99/30274

300 including multithreaded code constructs is processed by a conventional compiler 302. The compiler communicates with a linker 304 which links pre-existing routines from a library 306 with the output of the compiler to create
5 an executable module 308.

Existing operating systems, including the WIN32 API, often provide a general purpose thread library which may allow carrying out defined tasks like these. For example, a first thread may be defined for operating the CD ROM,
10 and another for the modem.

Windows NT WIN32 thread creation is unstructured. A thread is created by passing a function pointer and an argument pointer to a CREATETHREAD call. The new thread then executes the given function with the given argument.
15 There is no specific relationship between the created thread and the creating thread: the two threads are effectively asynchronous. One thread for example, can arbitrarily suspend, resume or terminate the execution of another thread.

20 This is not a problem for unrelated tasks like CD/modem tasks noted above. However, when two parts of a program are to be executed as threads, the synchronization operations are often complex and error prone. Unpredictable interactions among the multiple

WO 00/36491

PCT/US99/30274

threads can induce problems including race conditions, and deadlock. Effectively, the user is left with the daunting task of using these thread libraries in a way that does not cause this problem.

5 The present application discloses a specific embodiment operating using the Windows NT (TM) system. It should be understood, however, that this system is portable across many platforms and that the same concept described herein can be used in those systems, including
10 Linux, and any other operating system.

While a process has its own address space, a thread is often simply a program counter and stack pointer. A process may have many threads but all the threads share the same address space.

15 Figures 2 and 3 show this operation in a computer system. Figure 2 shows a computer system, with four processors 200, 202, 204, 206. The processors can be in a multiple processor system as shown. The pool 199 of processors is associated with an operating system 210, a
20 user interface 215, auxiliary hardware 220 (e.g. memory, chipsets, etc), a display 225 and other computer components.

The operating system 210 includes multiple threads 212, 214, and others. Each thread is resident on the

WO 00/36491

PCT/US99/30274

stack within the heap. The threads are associated with processors, which execute the threads. Figure 2 shows the pool of threads on the left and the pool of processors on the right. Each of the dynamically-created
5 threads are peers. Of course there can be many more threads than processors. The operating system controls the threads to dynamically switch between the processors.

The present application defines an entirely new way of creating, synchronizing, and handling the
10 synchronization among threads. The system uses a new way of compiling code based on multithreadable code, either alone, or in conjunction with multithreaded code.

The operating system or programming language has a higher level system that uses special constructs called
15 "equivalency annotations". A lower level function call based system is used with special objects. Those special objects can synchronize among the threads in a way that prevents the objects within the threads from having ambiguous states.

20 Many of these systems are based on the concept of equivalence annotations. Equivalence annotations can take many forms - pragmas, special keywords, special kinds of comments, special characters, textual modifications (such as boldfacing, underlining, or

WO 00/36491

PCT/US99/30274

italics), and others. They could be part of the program text, or in a separate file i.e., an extra file that contains nothing but the annotations. The pragma can have meaning to a compiler. Pragmas often form a
5 specified syntax, but usually convey nonessential information that is intended to help the compiler to optimize the program.

The present embodiment uses these pragmas as special equivalence annotations. Pragmas are convenient for
10 annotations since many programming language already provide pragmas for other purposes. While a pragma is described as being used as the preferred annotation of the present application, the program can certainly be annotated in other manners. For example, Java does not
15 support pragmas, so a special kind of comment line could be used. The equivalence annotations described throughout this specification should be understood to be interpretable in this way.

The MULTITHREADABLE equivalence annotation can be a
20 pragma when embodied in the C programming language. This indicates that a block or loop can be executed in a multithreaded manner. This means that there is no timing dependent nondeterminacy, and the system can execute the instructions into a multithreaded system.

WO 00/36491

PCT/US99/30274

The MULTITHREADED equivalence annotation means that a block or loop must be executed in a multithreaded manner. The multithreaded execution need not be equivalent to sequential execution. Lock synchronization can be used
5 to introduce nondeterminacy if desired.

The equivalence annotation becomes part of the operating system. Special, monotonically increasing and otherwise constrained s-COUNTERS, and similarly constrained s-FLAGS are operated to synchronize the access of threads
10 to shared memory in order to prevent unwanted interference.

A special SYNCHRONIZATION COUNTER, or s-COUNTER is defined as an object with three basic attributes. The s-COUNTER has a non-negative integer value. The object only allows
15 an INCREMENT operation and a CHECK operation. An initial value of the s-COUNTER object is set to zero. An INCREMENT function automatically increases the value of the counter by a specified amount. The CHECK operation suspends the calling thread until the value of the counter becomes
20 greater than or equal to a specified level.

The multi-threaded programming system has a higher level notation includes annotation objects in the program code. Using the example of the c language, this can be described as "multi-threaded c".

WO 00/36491

PCT/US99/30274

A lower level structured thread library is described as "Sthreads". The special annotation objects are transformed into special Sthread calls by the Sthreads annotation objects pre-processor.

5 The multithreaded model uses the thread synchronization/annotation objects disclosed above to synchronize among the threads.

Threads can be created in different ways. A first thread creation construct is the multithreaded block.

10 This is indicated by the MULTITHREADED keyword placed immediately before an ordinary C block:

```
15       MULTITHREADED {  
          statement  
          ...  
          statement  
      }
```

This notation specifies that the statements of the block should be executed as asynchronous threads. This

20 is a conventional way of referring to these threads. For example, the operating system could create threads to read from CD, and threads to read from tape. The threads are executed and proceed concurrently. They all share the same address space as the parent program. Execution

25 does not continue past the multithreaded block until all the threads have individually terminated. It is typically illegal for the program to contain any kind of jump

WO 00/36491

PCT/US99/30274

between the individual statements of the block, from inside the block to outside the block, or from outside the block to inside the block.

A second thread creation construct is the
5 multithreaded for-loop, indicated by the `MULTITHREADED`
keyword placed immediately before an ordinary for-loop:

```
10      MULTITHREADED  
      FOR (I = expression; I comparison expression; I = I +  
          expression)  
          statement
```

This notation specifies that the iterations of the loop should be executed as asynchronous threads. The threads all share the same address space as the parent
15 program. Each thread, however, has a local copy of the loop control-variable with a different value from the iteration range. The iteration scheme can restrict to a single control-variable and expressions that are not modified within the loop body. Execution does not
20 continue past the multithreaded for-loop until all the threads have individually terminated. It is illegal for the program to contain any kind of jump from inside the loop to outside the loop or from outside the loop to inside the loop. In essence, a multithreaded for-loop is
25 a quantified form of multithreaded block.

Multithreaded and ordinary blocks and for-loops can be arbitrarily nested.

WO 00/36491

PCT/US99/30274

Traditional approaches have often been categorized as either being explicitly multithreaded or implicitly multithreaded. With explicit multithreading, the programmer expresses exactly how the operations of the program are executed by threads. Implicit multithreading is carried out when the programmer writes an ordinary sequential program. The programming system, e.g. the compiler, determines how the operations can be executed by separate threads.

10 The present application goes beyond the multithreaded concepts described above into a concept of multithreadable code constructs. The multithreadable construct can be executed according to a specified sequential operational semantics. The most common

15 operational semantics would be executed sequentially. An alternative, however, allows the multithreadable code construct to be operated according to multithreaded operational semantics.

Rules are defined that constrain the multithreaded

20 execution such that its result is equivalent to sequential execution.

As disclosed herein, the multithreadable code construct is formed of:

WO 00/36491

PCT/US99/30274

- i) a syntactic description of the form of the construct,
- (ii) a sequential operational semantics, that, when executed, defines how to execute the construct
- 5 sequentially,
- (iii) a multithreaded operational semantics, defining how to execute the construct by a set of threads, and
- (iv) a set of implicit or explicit programming rules
- 10 that are sufficient to ensure the equivalence of sequential and multithreaded execution of the construct.

The MULTITHREADABLE pragma becomes an assertion by the programmer that the BLOCK or FOR loop can be executed in a multithreaded manner without changing the results of the

15 program. The MULTITHREADABLE pragma can be applied to blocks and for loops in which the statements or iterations are independent of each other. The multithreaded execution is equivalent to sequential execution in this case. It is not a directive that the BLOCK or FOR LOOP *must* be executed

20 in a multithreaded manner.

As a simple example, consider the following program to sum the elements of a two-dimensional array:

```
void SumElements(float A[N][N], float *sum, int numThreads)
{
25     int i;
        float rowSum[N];
```


WO 00/36491

PCT/US99/30274

```

        #pragma multithreadable mapping(blocked(numThreads))
        for (i = 0; i < N; i++) {
            int j;
5           rowSum[i] = 0.0;
            for (j = 0; j < N; j++)
                rowSum[i] = rowSum[i] + A[i][j];
        }
        *sum = 0.0;
10        for (i = 0; i < N; i++)
            *sum = *sum + rowSum[i];
    }

```

Multithreaded execution of the FOR loop is equivalent to sequential execution because the iterations all modify different ROWSUM[I] and J variables. The arguments following the pragma indicate that multithreaded execution should assign iterations to NUMTHREADS different threads using a blocked mapping. There is a rich set of options that control the mapping of iterations to threads.

Therefore, the Multithreaded C preprocessor has two modes: sequential mode in which the MULTITHREADABLE pragma is ignored, and multithreaded mode in which the MULTITHREADABLE pragma is transformed into Sthreads calls.

Programs can be developed, tested, and debugged in sequential mode, then executed in multithreaded mode for performance. In addition, performance analysis and tuning can often be performed in sequential mode.

Determinacy of results is an important consequence of the equivalence of multithreaded and sequential execution. If sequential execution is deterministic

WO 00/36491

PCT/US99/30274

(which is usually the case), multithreaded execution must also be deterministic. Determinacy is usually desirable, since program development and debugging can be difficult when different runs produce different results. In other

5 multithreaded programming systems, determinacy is difficult to ensure. For example, locks, semaphores, and many-to-one message passing almost always introduce race conditions and hence nondeterminacy. However, nondeterminacy is important for efficiency in some

10 algorithms, e.g., branch-and-bound algorithms.

Multithreaded and multithreadable code constructs are integrated in this system. The programming system incorporates both explicitly multithreaded constructs which must be executed according to the multithreaded

15 semantics, along with multithreadable constructs which may selectively executed according to their sequential or multithreaded semantics. The multithreaded constructs are generally used to express multithreaded algorithms that have no sequential equivalent. This can include

20 controlling different hardware that haw no integration with one another, or controlling simultaneous different windows in a graphical user interface.

Multithreadable constructs are used to express the opportunity to use multithreading to speed up the

WO 00/36491

PCT/US99/30274

execution of a computationally-intensive algorithm, by using multiple threads on multiple processors.

By using both multithreaded and multithreadable constructs, the operating system can use one thread to control each window with multithreaded constructs, and the output to each window, within a window, is computed faster with the multiple threads using multithreadable constructs.

As described above, the synchronization can be carried out by an entry S-COUNTER, or an S-FLAG. Each are defined to have certain constraints.

An S-COUNTER, defined in the context of the C programming language, is diagrammed in Figure 3. It can be defined as a type definition and a set of interface functions. The counters are encapsulated as a class in an object-oriented language such as C++ or Java. The definition of the fundamental programming interface for S-COUNTERS is as follows:

```

20     typedef counter type definition Counter;
        int InitializeCounter(Counter *c);
        /* Initializes value(c) to zero.
        */
25     /* Must be called only once, before all other operations on
        counter c. */

        int FinalizeCounter(Counter *c);
        /* Must be called only once, after all other operations on c. */

30     int CheckCounter(Counter *c, unsigned int level);
        /* Suspends until value(c) greaterorequal level. */

        int IncrementCounter(Counter *c, unsigned int amount);
        /* Increases value(c) by amount. */

```

WO 00/36491

PCT/US99/30274

An s-COUNTER object *c* implicitly has a nonnegative integer attribute *value(c)*, which can only be accessed through the interface functions. The INITIALIZE function at
5 300 initializes *value(c)* to zero or some initial value.

Importantly, the counter is monotonic, as illustrated in 310. No decrement function is defined. Its value monotonically increases.

Any attempt to CHECK the counter, shown as step 320,
10 suspends the calling thread at 325. This prevents a condition which can catch or miss some action occurring during the check operation. Each s-COUNTER maintains a dynamic list of thread suspension queues 330, with one queue for each value on which at least one Check
15 operation is suspended.

CHECK compares *value(c)* to *level* and suspends until *value(c)* becomes greater than or equal to *level*. This is generically shown as AWAKE in step 340. Increment at 310 atomically increases *value(c)* by *amount*, thereby
20 reawakening all Check operations suspended on values less than or equal to the new *value(c)*.

All the functions can return an error code. Possible error conditions include invalid arguments, operations on an uninitialized counter, and counter overflow.

WO 00/36491

PCT/US99/30274

The type definition described above is carefully selected to remove the possibility of race conditions occurring on counter synchronization. There is no DECREASE operation. Therefore, the value of an s-COUNTER is

5 monotonically increasing. There is no possibility of a CHECK operation missing an INCREMENT operation since check suspends the thread. There is no PROBE or nonblocking CHECK operation. It is recognized by the inventor that any instantaneous value may depend on the relative timing of

10 the individual threads. Therefore, no operation can be based on the instantaneous value of an s-COUNTER.

A RESET operation can also be used to efficiently reuse counters between different phases of a program.

Alternatively, the old counters can be deleted and

15 recreated as new counters. RESET simply resets value(c) back to zero. However, to avoid the possibility of race conditions, RESET must not be called concurrently with any other operation on the same counter. RESET ends the process, and is not intended as a means of

20 synchronization between threads.

Another thread synchronization object is a special flag, called an s-FLAG. S-FLAGS, like s-COUNTERS, have restricted allowed operations within the multithreadable code concept. S-FLAGS support SET and CHECK operations.

WO 00/36491

PCT/US99/30274

Initially, an s-FLAG is not set. A SET operation on an s-FLAG atomically sets the flag. A CHECK operation on a flag suspends until the flag is set. Once an s-FLAG is set, it remains set.

5 Flags and counters are provided to provide deterministic synchronization within multithreadable constructs, as previously described.

 In summary of the above, an s-COUNTER object has the following operations (expressed in the C programming
10 language):

```
Initialize(Counter *c)
Finalize(Counter *c)
Increment(Counter *c, unsigned int amount)
15   Check(Counter *c, unsigned int value)
Reset(Counter *c)
```

 The Initialize operation initializes the Counter object and sets its value to zero. The Finalize operation
20 destroys the Counter object. An Increment operation increases the value of the Counter object by amount. A Check operation suspends the calling thread until the value of the Counter object is at least value. A Reset operation resets the value of the Counter object to zero.

25 In the following simple example, a "producer thread" produces items and writes them to a buffer. A group of

WO 00/36491

PCT/US99/30274

one or more concurrently executed "consumer threads" each independently reads the items from the buffer. The key synchronization issue is to prevent the consumer threads from reading items from the buffer that have not yet been
5 written by the producer thread. The following program fragment gives implementations of the producer thread and consumer threads (in the C programming language) using a counter for synchronization.

```
10     Counter count;
        Item buffer[NUM_ITEMS];

        ProducerThread(int blockSize)
        {
15             int index = 0, c = 0;
                while (index < NUM_ITEMS) {
                    buffer[index] = Produce();
                    index = index + 1;
                    c = c + 1;
20                     if (c == blockSize) {
                        Increment(count, blockSize);
                        c = 0;
                    }
                }
25     }

        ConsumerThread(int blockSize)
        {
            int index = 0, c = blockSize;
30             while (index < NUM_ITEMS) {
                if (c == blockSize) {
                    Check(count, index + blockSize);
                    c = 0;
                }
35             Consume(buffer[index]);
                index = index + 1;
                c = c + 1;
            }
        }
```

WO 00/36491

PCT/US99/30274

After writing a block of items to the buffer, the producer thread increments the s-COUNTER. Before reading a block of items from the buffer, a consumer thread checks
5 the counter. If the next block of items has not yet been written to the buffer, the consumer thread suspends until enough items have been written. The program does not require that the producer and consumer threads all use the same blockSize values.

10 The monotonicity of counters helps guarantee deterministic synchronization and the equivalence of multithreaded and sequential execution.

If shared variables are guarded against concurrent operations, a program that uses only counter
15 synchronization can produce deterministic results on all executions. Moreover, if sequential execution of the program (i.e., execution ignoring the MULTITHREADED keyword) does not deadlock, multithreaded execution is guaranteed not to deadlock and to produce the same results as
20 sequential execution. These properties are extremely useful in the testing and debugging of multithreaded programs.

Even in the absence of concurrent operations on shared variables, traditional synchronization mechanisms
25 can introduce nondeterminacy into a program through

WO 00/36491

PCT/US99/30274

timing dependent race conditions between threads. For example, consider the following program that uses a lock:

```
5      multithreaded {  
        { AcquireLock(&xLock); x = x+1; ReleaseLock(&xLock); }  
        { AcquireLock(&xLock); x = x*2; ReleaseLock(&xLock); }  
      }
```

Even though there are no concurrent operations on x, the resulting value of x is nondeterministic because of the race condition on the order in which the two threads acquire the lock. In contrast, because counters are monotonic, once a synchronization condition is enabled it remains enabled, and there is no possibility of a race condition to catch or miss a particular counter value.

15 For example, consider the following program that uses a counter:

```
20      multithreaded {  
        { CheckCounter(&xCount, 0); x = x+1;  
          IncrementCounter(&xCount, 1); }  
        { CheckCounter(&xCount, 1); x = x*2;  
          IncrementCounter(&xCount, 1); }  
      }
```

The resulting value of x is deterministic, because the CHECKCOUNTER operations will succeed in the same order in all executions, therefore the operations on x will occur in the same order. Moreover, since sequential execution does not deadlock, multithreaded execution cannot deadlock and will always produce the same results as sequential execution.

30

WO 00/36491

PCT/US99/30274

Programs that use only counter synchronization can still be erroneously nondeterministic if they do not guard against concurrent access to shared variables. For example, consider the following program using a counter:

```

5      multithreaded {
        { CheckCounter(&xCount, 0); x = x+1;
          IncrementCounter(&xCount, 1); }
        { CheckCounter(&xCount, 0); x = x*2;
          IncrementCounter(&xCount, 1); }
10     }

```

The result of the program is nondeterministic because of the possibility of concurrent execution of the operations on x. The nondeterminacy is caused by concurrent access to a shared variable, not by a synchronization race condition.

As a simple example, consider the following program to sum the elements of a two-dimensional array:

```

20     void SumElements(float A[N][N], float *sum, int numThreads)
        {
            int i;
            SthreadCounter counter;

            SthreadCounterInitialize(&counter);
25     #pragma multithreadable mapping(blocked(numThreads))
            for (i = 0; i < N; i++) {
                int j;
                float rowSum;
                rowSum = 0.0;
30     for (j = 0; j < N; j++)
                    rowSum = rowSum + A[i][j];
                SthreadCounterCheck(&counter, i);
                *sum = *sum + rowSum;
                SthreadCounterIncrement(&counter, 1);
35     }
            SthreadCounterFinalize(&counter);
        }

```

WO 00/36491

PCT/US99/30274

Without the counter operations, multithreaded execution of the for loop would not be equivalent to sequential execution, because the iterations all modify the same *sum variable. However, the counter operations ensure that multithreaded execution is equivalent to sequential execution. In sequential execution, the iterations are executed in increasing order and the STHREADCOUNTERCHECK operations succeed without suspending. In multithreaded execution, the counter operations ensure that the operations on *sum occur atomically and in the same order as in sequential execution. Iteration i suspends at the STHREADCOUNTERCHECK operation until iteration i - 1 has executed the STHREADCOUNTERINCREMENT operation.

Conditions are carved out to prevent concurrent access to shared variables using counters. Essentially, each pair of operations on a shared variable must be separated by a transitive chain of counter operations. If these conditions can be shown to hold in any one execution of the program, they must hold in all executions of the program. Therefore, if sequential execution satisfies the conditions, multithreaded execution is also guaranteed to satisfy the conditions, hence produce the same results as sequential execution. This result forms the basis of a powerful methodology for

WO 00/36491

PCT/US99/30274

developing multithreaded programs using sequential reasoning, testing, and debugging techniques.

All the programs using counters so far discussed satisfy the conditions on shared variables, therefore are
5 guaranteed to be deterministic. In addition, the program examples described herein have equivalent multithreaded and sequential execution. The cost of increased determinacy is decreased concurrency. Synchronization using counters provides an effective means of controlling
10 this tradeoff between determinacy and concurrency.

Counters can also be used as a stronger form of lock synchronization, providing sequential ordering in addition to mutual exclusion on a critical section. With the traditional implementation of mutual exclusion using
15 a pair of lock operations, the order in which threads enter the critical section is nondeterministic. This is desirable in terms of maximizing concurrency, but is undesirable in terms of reasoning, testing, and debugging, and simply might not satisfy the desired
20 program specification. Replacing the pair of lock operations with a pair of counter operations can guarantee deterministic results, at the cost of decreased opportunities for concurrency.

WO 00/36491

PCT/US99/30274

Consider the computation of a result object formed by accumulating a series of independent subresults that are computed concurrently. For example, the result object could be a linked list and the accumulate operation could be an append, or the result object could be a summation and the accumulate operation could be an addition. Mutual exclusion is required to prevent interference between multiple concurrent accumulate operations on the result object.

10 The following program implements the computation with one thread computing each subresult, and a pair of lock operations to provide mutual exclusion:

```
15       CompositeItem result;  
      Lock resultLock;  
      InitializeLock(&resultLock);  
      multithreaded for (i = 0; i < N; i++) {  
          SingleItem subresult = compute(i);  
          AcquireLock(&resultLock);  
20       accumulate(&result, subresult);  
          ReleaseLock(&resultLock);  
      }  
      FinalizeLock(&resultLock);
```

25 Only one thread can hold RESULTLOCK at any given time, thereby ensuring mutual exclusion of the ACCUMULATE operations. However, if the accumulate operation is not associative and determinacy of results is desired, some other mechanism is required to ensure sequential (or at least deterministic) ordering, in addition to mutual exclusion. For example, neither appending an item to a linked list, nor floating point addition are associative

WO 00/36491

PCT/US99/30274

operations. With both these examples, the above program may produce different results on repeated executions.

The following program implements the computation with the pair of lock operations replaced with a pair of
 5 counter operations, to provide both mutual exclusion and sequential ordering:

```

    CompositeItem result;
    Counter resultCount;
10  InitializeCounter(&resultCount);
    multithreaded for (i = 0; i < N; i++) {
        SingleItem subresult = compute(i);
        CheckCounter(&resultCount, i);
        accumulate(&result, subresult);
15  IncrementCounter(&resultCount, 1);
    }
    FinalizeCounter(&resultCount);
  
```

As with the lock program, only one ACCUMULATE
 20 operation can execute at any given time. However, the accumulate operations are now additionally constrained to execute in sequential order. RESULTCOUNT[I] = I indicates that thread I-1 has completed its accumulate operation. The counter program has greater determinacy at the cost
 25 of less concurrency. With the lock program, an ACCUMULATE operation can execute concurrently with compute operations in all other threads. With the counter program, an ACCUMULATE operation can only execute
 30 concurrently with compute operations in higher numbered threads.

The optimal tradeoff between determinacy and concurrency has to be made on a case by case basis.

WO 00/36491

PCT/US99/30274

Counters are a powerful mechanism for providing sequential ordering on top of mutual exclusion in the many cases where determinacy is important and the performance consequences of less concurrency are not
 5 great.

The Sthreads Interface

The code produced according to the present application can be expressed using the Multithreaded C
 10 pragma notation. As described previously, there is a direct correspondence between the pragma notation for thread creation and the Sthreads library functions that support thread creation. As a simple example, the following is a program implemented using Sthreads:

```

15     typedef struct {
        float (*A) [N];
        float *sum;
        SthreadCounter *counter;
    } LoopArgs;
20
    void LoopBody(int i, int notused1, int notused2, LoopArgs
    *args)
    {
        int j;
        float rowSum;
25         rowSum = 0.0;
        for (j = 0; j < N; j++)
            rowSum = rowSum + (args->A) [i] [j];
        SthreadCounterCheck(args->counter, i);
30         *(args->sum) = *(args->sum) + rowSum;
        SthreadCounterIncrement(args->counter, 1);
    }

    void SumElements(float A[N] [N], float *sum, int numThreads)
35     {
        int i;
        SthreadCounter counter;
  
```

WO 00/36491

PCT/US99/30274

```

    LoopArgs args;

    SthreadCounterInitialize(&counter);
    args.A = A;
5   args.sum = sum;
    args.counter = &counter;
    SthreadRegularForLoop(
        (void (*)(int, int, int, void *)) LoopBody,
        (void *) &LoopArgs,
10   0, STHREAD_CONDITION_LT, N, 1,
        1, STHREAD_MAPPING_BLOCKED, numThreads,
        STHREAD_PRIORITY_PARENT, STHREAD_STACK_SIZE_DEFAULT);
    SthreadCounterFinalize(&counter);
}

15
```

Although this program is syntactically more complicated than the Multithreaded C version, it is considerably less complicated than the same program expressed using Windows NT threads. The mechanics of creating threads, assigning iterations to threads, and waiting for thread termination is handled within the Sthreads library call.

The Sthreads multithreaded programming system is implemented as a transparent add-on to an existing program development system, e.g., a compiler or interpreter, or other program development environment. The notation and implementation allows multithreaded and multithreadable code constructs to be directly translated into a high-level structured thread library. This translation is implemented as a preprocessor that can be transparently called prior to the standard compilation phase in an existing program development system.

WO 00/36491

PCT/US99/30274

For example, when integrated with the Microsoft Visual C++ programming system, the standard CL (Compiler-Linker) is replaced by a special Sthreads tool that calls the Sthreads preprocessor on program files, then calls
5 the standard (renamed) Visual C++ CL.

Integration of Sthreads with existing programming systems allows programmers new flexibility without adopting new programming systems to use the power of multithreading. They can use their standard editor,
10 debugger, compiler, etc., and simply add Sthreads to the system. It also means that Sthreads piggybacks on the quality of code generation and error analysis of the underlying development system.

Preprocessing had previously been used for many
15 kinds of program "source-to-source" transformations. Sthreads in contrast, implements a full-fledged, sophisticated multithreaded programming system by using a preprocessing integrated with a standard program development environment.

20 One implementation has been created in the ANSI C language, thereby defining a "Multithreaded C" language. A structured thread library (Sthreads) is called by the languages. In both Multithreaded C and Sthreads, thread creation constructs are multithreaded variants of

WO 00/36491

PCT/US99/30274

sequential "block" (i.e., sections of program code distinctly defined by conventional program constructs) and "for loop" constructs. In the Multithreaded C implementation, these constructs are supported as pragma annotations to a sequential program. With Sthreads, exactly the same constructs are supported as library calls. At both levels, synchronization objects and operations are supported as Sthreads library calls. In this embodiment, the Sthreads library for Windows NT can be implemented as a very thin layer on top of the Win32 thread API. As a consequence, there is essentially no performance overhead associated with using Sthreads or Multithreadable C, as compared to using Win32 threads directly.

15 Multithreaded C is implemented as a portable source-to-source preprocessor that directly transforms annotated blocks and for loops into equivalent calls to the Sthreads library. The programmer has the option of either using the pragma annotations and preprocessor or making
20 Sthreads calls directly. The Sthreads library and Multithreaded C preprocessor are integrated with Microsoft Developer Studio Visual C++. Building a project preferably automatically invokes the preprocessor where necessary and links with the Sthreads library.

WO 00/36491

PCT/US99/30274

Multithreadable blocks and for loops are implemented as a sequence of `CREATETHREAD` calls followed by a `WAITFOR-SINGLEOBJECT` call on an event. Terminating threads perform an `INTERLOCKEDDECREMENT` call on a counter, and the last
5 thread to terminate performs a `SETEVENT` call on the event. Flags are implemented directly as Win32 events. Counters are implemented as linked lists of Win32 events, with one event for every value on which some thread is waiting. Locks are implemented directly as Win32 critical
10 sections. Barriers are implemented as a pair of Win32 events and a Win32 critical section.

An important issue of the multithreading operation comes about when considering multiple processors. The hardware and operating systems of modern technology allow
15 for multiprocessor systems. Current operation in multiprocessor systems, however, have often simply operated on one but not the other processor. By multithreading in this way, the different threads can actually be executed on the different processors.

20 In operation, when a multithreading indicator (such as a "compile as multithreaded" flag/button/environment-variable) is set, both multithreadable and multithreaded blocks/loops are compiled to multithreaded code. When the multithreading indicator is not set, the

WO 00/36491

PCT/US99/30274

multithreaded blocks/loops are compiled to multithreaded code and the multithreadable blocks/loops are compiled into ordinary sequential code. This allows a programmer to mix constructs that only have a multithreaded meaning
5 (e.g., real-time control and systems programming uses of threads) with constructs that can be compiled into threads for multiprocessor performance or compiled into equivalent sequential code when developing and debugging.

The invention allows a program to run as fast as a
10 sequential program on one processor, but significantly faster on multiprocessors, without recompilation, relinking, or reconfiguration. The invention thus allows a program to adapt dynamically to changing resources. Use of monotonic flags and monotonic counters makes
15 embodiments of the invention reliable and timely. The mapping of Statements/Iterations onto threads is relatively simple. One thread is used for each statement/chunk, or for a small number of statements/chunks.

20 A Typical for loop may have thousands or millions of iterations. The Overhead associated with assigning units of work to threads is significant. The present application defines assigning the iterations in

WO 00/36491

PCT/US99/30274

contiguous "chunks". Significant unit of work should be performed by each chunk.

For example:

```
#PRAGMA MULTITHREADABLE CHUNKSIZE(1000), MAPPING(BLOCKED(T))  
5   FOR (I = 0; I < N; I++)  
      A[I] = B[I];
```

Interaction of Chunksize and Mapping is described in the following example:

```
10  
    #PRAGMA MULTITHREADABLE CHUNKSIZE(2), MAPPING(BLOCKED(4))  
    FOR (I = 50; I >= 10; I = I - 2)  
        DoSOMETHING(I);
```

15 The Complete Sthreads Library includes a number of statements:

Processor Management: STHREADSGETNUMPROCESSORSPRESENT,
STHEADSSETNUMPROCESSORSUSED, STHREADSGETPROCESSORSPRESENT,
STHEADSSETPROCESSORSUSED.

20 Thread Creation: STHREADSBLOCK, STHREADSREGULARFORLOOP.

Thread Scheduling: STHREADSGETCURRENTPRIORITY,
STHEADSSETCURRENTPRIORITY.

Flags: STHREADSFLAGINITIALIZE, STHREADSFLAGFINALIZE,
STHEADSFLAGSET, STHREADSFLAGCHECK, STHREADSFLAGRESET.

WO 00/36491

PCT/US99/30274

Counters: STHREADSCOUNTERINITIALIZE, STHREADSCOUNTERFINALIZE,
STTHREADSCOUNTERINCREMENT, STHREADSCOUNTERCHECK, STHREADSCOUNTERRESET.

Locks: STHREADSLOCKINITIALIZE, STHREADSLOCKFINALIZE,
STTHREADSLOCKACQUIRE, STHREADSLOCKRELEASE.

5 Barriers: STHREADSBARRIERINITIALIZE, STHREADSBARRIERFINALIZE,
STTHREADSBARRIERPASS, STHREADSBARRIERRESET.

Examples of computer program code implementing each
of these constructs are set forth in the appendix.

FIG. 4 is a process flowchart showing a method for
10 compiling multithreadable code in accordance with one
embodiment of the invention. The computer program source
code text 400 includes annotations defining
multithreadable code constructs (and, optionally,
multithreaded code constructs) and any necessary
15 processor management, thread creation, and
synchronization constructs (such as monotonic flags and
counters). If a multithreading indicator is set 401, the
source code text 400 is processed by a pre-processor 402
that parses the source into an expanded computer program
20 text. The expanded computer program text includes
inserted calls to an Sthreads library 406 to invoke
multithreaded program operations wherever a source code
annotation called for multithreadable functionality. A
conventional compiler 406 then communicates with a linker

WO 00/36491

PCT/US99/30274

410 which links pre-existing routines from the Sthreads library 406 with the output of the compiler to create an executable module 412.

If the multithreading indicator is not set 401, the
5 original computer program source code text 400 is compiled and linked in conventional fashion, with each section of multithreadable code constructs compiled as sequentially executing code. Annotations not recognized by the compiler 408 are ignored.

10 A convenient implementation shortcut that permits ready use of conventional compilers and linkers is to rename a pre-existing compiler-linker executable file to a new name, and assign the old name of the compiler-linker executable file to the pre-processor. The pre-processor then can call the compiler-linker executable
15 file when needed by invoking the new name.

Synchronization Using Locks

Locks are provided to express nondeterministic
20 synchronization, usually mutual exclusion, within multithreaded BLOCKS and FOR loops. Sthread locks support the usual ACQUIRE and RELEASE operations. The order in which concurrent ACQUIRE operations succeed is nondeterministic. Therefore, there is very little use for locks within

WO 00/36491

PCT/US99/30274

multithreadable blocks and FOR loops. As a simple example, consider the following program to sum the elements of a two-dimensional array:

```

5      void SumElements(float A[N][N], float *sum, int numThreads)
      {
          int i;
          SthreadLock lock;

          SthreadLockInitialize(&lock);
10      #pragma multithreaded mapping(blocked(numThreads))
          for (i = 0; i < N; i++) {
              int j;
              float rowSum;
              rowSum = 0.0;
15      for (j = 0; j < N; j++)
                  rowSum = rowSum + A[i][j];
                  SthreadLockAcquire(&lock);
                  *sum = *sum + rowSum;
20      SthreadLockRelease(&lock);
          }
          SthreadLockFinalize(&lock);
      }

```

Like the flag operations in the program, the lock

25 operations in this program ensure that the operations on *sum occur atomically. However, unlike the flag operations, the lock operations do not ensure that the operations on *sum occur in the same order as in sequential execution, or even in the same order each time

30 the program is executed. Therefore, since floating-point addition is not associative, the program may produce different results each time it is executed. However, because execution order is less restricted, this program allows more concurrency than the program described above.

WO 00/36491

PCT/US99/30274

This is an example of the commonly-occurring tradeoff between determinacy and efficiency.

Synchronization Using Barriers

5 S-thread barriers are provided to express collective
synchronization of a group of threads in cases when
thread termination and recreation is too expensive. The
barriers described herein support the usual Pass
operation. All the threads in a group must enter the Pass
10 operation before all the threads in the group are allowed
to leave the Pass operation. In current systems, the cost
of N threads executing a Pass operation is less than the
cost of creating and terminating N threads. Therefore, a
typical use of barriers is to replace a sequence of
15 multithreadable loops with a single multithreaded loop
containing a sequence of barrier Pass operations. However,
with modern lightweight thread systems such as
Windows NT, we are discovering that barriers are required
for efficiency in very few circumstances.

20 A number of examples are described herein.

Trivial Example: Independent Iterations

```
INT ARRAYSUM(FLOAT A[N] [N])  
/* SUMS THE ELEMENTS OF A 2-DIMENSIONAL ARRAY. */  
{  
25       FLOAT SUM, ROWSUM[N];  
      INT I;
```

WO 00/36491

PCT/US99/30274

```

SUM = 0.0;
#pragma MULTITHREADABLE
FOR (I = 0; I < N; I++) {
    INT J;
5    ROWSUM[I] = 0.0;
    FOR (J = 0; J < N; J++)
        ROWSUM[I] = ROWSUM[I] + A[I][J];
}
10 FOR (I = 0; I < N; I++)
    SUM = SUM + ROWSUM[I];
RETURN SUM;
}

```

A more difficult example is shown in the following.

```

15
INCORRECT EXAMPLE: NONDETERMINACY
INT ARRAYSUM(FLOAT A[N][N])
/* SUMS THE ELEMENTS OF A 2-DIMENSIONAL ARRAY. */
{
20    FLOAT SUM;
    INT I;
    STTHREADSLOCK SUMLOCK;

    STTHREADSLOCKINITIALIZE(&SUMLOCK);
25    SUM = 0.0;
    #pragma MULTITHREADABLE
    FOR (I = 0; I < N; I++) {
        INT J;
        FLOAT ROWSUM = 0.0;
30        FOR (J = 0; J < N; J++)
            ROWSUM = ROWSUM + A[I][J];
        STTHREADSLOCKACQUIRE(&SUMLOCK);
        SUM = SUM + ROWSUM;
        STTHREADSLOCKRELEASE(&SUMLOCK);
35    }
    STTHREADSLOCKFINALIZE(&SUMLOCK);
    RETURN SUM;
}
INT ARRAYSUM(FLOAT A[N][N])
40 /* SUMS THE ELEMENTS OF A 2-DIMENSIONAL ARRAY. */
{
    FLOAT SUM;
    INT I;
    STTHREADSCOUNTER SUMCOUNT;
45

```

WO 00/36491

PCT/US99/30274

```

    STHREADSCOUNTERINITIALIZE(&SUMCOUNT);
    SUM = 0.0;
    #PRAGMA MULTITHREADABLE
    FOR (I = 0; I < N; I++) {
5       INT J;
        FLOAT ROWSUM = 0.0;
        FOR (J = 0; J < N; J++)
            ROWSUM = ROWSUM + A[I][J];
        STHREADSCOUNTERCHECK(&SUMCOUNT, I);
10      SUM = SUM + ROWSUM;
        STHREADSCOUNTERINCREMENT(&SUMCOUNT, 1);
    }
    STHREADSCOUNTERFINALIZE(&SUMCOUNT);
    RETURN SUM;
15

```

As can be seen, iterations cannot be executed as separate threads because of nondeterminacy in the top. However, the counters allow determinacy between the system therefore enabling the system to be multithreaded.

20

Single-Writer Multiple-Reader Broadcast

Counters can be used to provide elegant, flexible, and efficient dataflow synchronization between a single writer and multiple readers of a sequence of items written to an array. In this synchronization pattern, reading an item does not remove it from the sequence—each reader independently reads the entire shared array. Because a counter has multiple thread suspension queues, a single counter object can be used to synchronize the writer thread and any number of completely independent reader threads, with each thread potentially having a different granularity of synchronization. The writer

30

WO 00/36491

PCT/US99/30274

thread incrementing the counter broadcasts the availability of data to the entire set of reader threads.

The following program demonstrates the single-write multiple-reader broadcast pattern with synchronization on

5 every item:

```

void Writer(Item *data, int n, Counter *dataCount)
{
    int i;
    for (i = 0; i < n; i++) {
10      data[i] = GenerateItem(i);
        IncrementCounter(dataCount, 1);
    }
}

void Reader(Item *data, int n, Counter *dataCount)
15 {
    int i;
    for (i = 0; i < n; i++) {
        CheckCounter(dataCount, i+1);
20      UseItem(data[i]);
    }
}

...
Item data[N];
Counter dataCount;
25 int r;

InitializeCounter(&dataCount);
multithreaded {
    Writer(data, N, dataCount);
30    multithreaded for (r = 0; r < numReaders; r++)
        Reader(data, N, dataCount);
}
FinalizeCounter(&dataCount);
...
35
```

One WRITER thread and an arbitrary number of Reader threads are executed concurrently, with communication through the shared data array, and synchronization through the DATACOUNT counter. At any point, some Reader

40 threads may be suspended in their CHECKCOUNTER operation, waiting for the Writer thread to increment DATACOUNT, while other Reader threads may be reading data items that have previously been written. The READER threads execute

WO 00/36491

PCT/US99/30274

independently of each other and do not synchronize their actions in any manner. The synchronization pattern is strictly a one-to-many broadcast from the WRITER thread to the READER threads.

5 Synchronization on every item that is written and read may be too expensive if the time taken to generate and use an item is too small. The single-reader multiple-writer broadcast pattern can be generalized to allow the writer and each reader thread to synchronize on a block
10 of items instead of on individual items. The following program adds an individual granularity of blocked synchronization to the writer and each reader thread:

```

void Writer(Item *data, int n, Counter dataCount, int blockSize)
{
    int i;
15     for (i = 0; i < n; i++) {
        data[i] = GenerateItem(i);
        if ((i+1)%blockSize == 0)
            IncrementCounter(dataCount, blockSize);
    }
20     IncrementCounter(dataCount, n-(n/blockSize)blockSize);
}

void Reader(Item *data, int n, Counter *dataCount, int
blockSize)
{
25     int i;
    for (i = 0; i < n; i++) {
        if (i%blockSize == 0)
            CheckCounter(dataCount, min(i+blockSize, n));
        UseItem(data[i]);
30     }
}

```

The WRITER and READER threads now increment and check the DATACOUNT counter in multiples of BLOCKSIZE and write and
35 read the data array in blocks of items. There is no requirement that BLOCKSIZE be the same in all threads. Different threads can be passed different BLOCKSIZE based

WO 00/36491

PCT/US99/30274

on their individual performance characteristics and requirements. This pattern is now extremely flexible and easily adaptable with regard to practical performance tuning.

- 5 The single-writer multiple-reader broadcast pattern is a dataflow synchronization pattern that occurs in many diverse applications of threads to multiprocessing. For example, in the Paraffins Problem, an array of molecules of a certain size can be generated by one thread and
- 10 concurrently read by other threads that in turn generate arrays of larger molecules. The pattern is very different from, for instance, the multiple-writers multiple-readers bounded-buffer problem, which is elegantly solved using semaphores. Just as counters are not well suited to
- 15 implementing bounded buffers, semaphores and other traditional synchronization mechanisms are not well suited to implementing the single-writer multiple-reader broadcast pattern.

20 Another Example Application: Aircraft Route Optimization

 The Aircraft Route Optimization Problem is part of the U.S. Air Force Rome Laboratory C3I Parallel Benchmark Suite. For this application, we achieved better performance using Sthreads on a quad-processor Pentium

WO 00/36491

PCT/US99/30274

Pro system running Windows NT than the best reported results for message-passing programs running on expensive Cray and SGI supercomputers with up to 64 processors. The flexibility of shared-memory, lightweight multithreading, and sequential development methods allowed us to develop a much more sophisticated and efficient algorithm than would be possible on a message-passing supercomputer.

The C3I Parallel Benchmark Suite

The U.S. Air Force Rome Laboratory C3I Parallel Benchmark Suite consists of eight problems chosen to represent the essential elements of real C3I (Command, Control, Communication, and Intelligence) applications. Each problem consists of the following:

A problem description giving the inputs and required outputs.

An efficient sequential program (written in C) to solve the problem.

The benchmark input data.

A correctness test for the benchmark output data.

For some of the problems, a parallel message-passing program is also provided. Rome Laboratory maintains a publicly accessible database of reported performance results.

WO 00/36491

PCT/US99/30274

The C3I Parallel Benchmark Suite provides a good framework for evaluating our structured multithreaded programming system. The problems are computationally intensive and involve a variety of complex algorithms and data structures. The sequential program provides us with a good starting point and a fair basis for performance comparison. The performance database allows us to compare our results with those of other researchers. For these reasons, we are developing multithreaded solutions to several of the C3I Parallel Benchmark Suite problems.

The task in the Aircraft Route Optimization Problem is to find the lowest-risk path for an aircraft from an origin point to a set of destination points in the airspace over an uneven terrain. The risk associated with each transition in the airspace is determined by its proximity to a set of threats. The problem involves realistic constraints on aircraft speed and maneuverability. The aircraft is also constrained to fly above the underlying terrain and beneath a given ceiling altitude.

The problem is essentially the single-source, multiple-destination shortest path problem with a large, sparsely connected graph. The airspace for the benchmark is 100 km by 100 km in area and 10 km in altitude,

WO 00/36491

PCT/US99/30274

discretized at 1 km intervals. The 100,000 positions in space correspond to 2,600,000 nodes in the graph, since each position can be reached from 26 different directions. Because of aircraft speed and maneuverability constraints, each node is connected to only nine or ten geographically adjacent nodes. Therefore, the graph consists of approximately 2.6 million nodes and 26 million edges.

The sequential algorithm to solve the Aircraft Route Optimization Problem is based on a queue of nodes. Initially the queue is empty except for the origin node. At each step, one node is removed from the queue. Valid transitions from this source node to all adjacent destination nodes are considered. For each destination node, if the path to the node via the source node is shorter than the current shortest path to the node, the path to the node is updated and the node added to the queue. The algorithm continues until the queue is empty, at which stage the shortest paths to all reachable nodes have been computed.

The queue is ordered on path length so that shorter paths are expanded before longer paths. This has a significant effect on performance. Without ordering, longer paths are expanded, then discarded when shorter

WO 00/36491

PCT/US99/30274

paths to the same points are expanded later in the computation. However, whether the queue is ordered, partially ordered, or unordered does not affect the results of the algorithm.

- 5 The most straightforward approach to obtaining parallelism in the Aircraft Route Optimization Problem is to geographically partition the airspace into blocks, with one thread (or process) responsible for each block. Each thread runs the sequential algorithm on its own
- 10 block using its own local queue and periodically exchanges boundary values with neighboring blocks. This approach is particularly appealing on distributed-memory, message-passing platforms, because memory can be permanently distributed according to the blocking
- 15 pattern. If the threads execute a reasonably large number of iterations between boundary exchanges, good load balance can be achieved.

- The problem with this algorithm is that, as the number of blocks/threads is increased the total amount of
- 20 computation also increases. Therefore, any speedup is based on an increasingly inefficient underlying algorithm. At any time, the local queues in most blocks contain paths that are too long and are irrelevant to the actual shortest paths. The processors are kept busy

WO 00/36491

PCT/US99/30274

performing computation that is later discarded. At any given time, it is only productive to work on an irregular and unpredictable subset of the graph. However, irregular and adaptive blocking schemes do not solve the problem,
5 since there is usually equal work available in all blocks. The issue is the distinction between productive and unproductive work.

Our solution is to statically partition the airspace into a large number of blocks and to use a much smaller
10 number of threads. A measure of the average path length is maintained with each local queue. At each step, the blocks with local queues containing the shortest paths are assigned to the threads. Therefore, the subset of blocks that are active and the assignment of blocks to
15 threads change dynamically throughout program execution. This algorithm takes advantage of the symmetric multiprocessing model, in which all threads can access the entire memory space with uniform cost. It also takes advantage of the lightweight multithreading model to
20 achieve good load balance, since the workload within each thread at each step is highly variable.

The ability to develop, test, and debug using sequential methods was crucial in the development of this sophisticated multithreaded algorithm. The entire program

WO 00/36491

PCT/US99/30274

was tested and debugged in sequential mode before
multithreaded execution was attempted. In particular,
development of the complex boundary exchange and queue
update algorithms would have been considerably more
5 difficult in multithreaded mode.

The ability to analyze and tune performance using
sequential methods was also very important. Good
performance depended on exposing enough parallelism
without significantly increasing the total amount of
10 computation. We determined efficient values for the
number of blocks, the number of threads, and the number
of iterations between boundary exchanges by measuring
computation times and operation counts of the
multithreaded program in running in sequential mode. This
15 detailed analysis would have been very difficult to
perform in multithreaded mode. We avoided memory
contention in multithreaded mode by avoiding cache misses
in sequential mode. The analysis of memory access
patterns in sequential mode is much simpler than in
20 multithreaded mode.

All Pairs Shortest Paths Example

This example describes the algorithmic and
performance advantages of counter synchronization. In the

WO 00/36491

PCT/US99/30274

example, a counter is used as a less restrictive, and consequently more efficient, replacement for a barrier. The example program is a multithreaded solution to the all-pairs shortest-path problem using the Floyd-Warshall
5 algorithm. Using traditional synchronization mechanisms, this problem can be solved using one barrier or, more efficiently, an array of condition variables. We show how the efficient solution can be implemented using a single counter instead of an array of condition variables. We
10 give timing measurements comparing the performance of the barrier, condition variable, and counter algorithms.

The all-pairs shortest-path problem takes as input the edge-weight matrix of a weighted directed graph, and returns the matrix of shortest-length paths between all
15 pairs of vertices in the graph. The graph is required to have no cycles of negative length, and the weight of the edge from a vertex to itself is required to be zero.

The following program solves the all-pairs shortest-path problem using the sequential Floyd-Warshall
20 algorithm:

WO 00/36491

PCT/US99/30274

```

    VOID SHORTESTPATHS1 (INT EDGE[N] [N] , INT PATH[N] [N])
    {
        INT K, I, J;
5        PATH[0..N-1][0..N-1] = EDGE[0..N-1][0..N-1];
        FOR (K = 0; K < N; K++)
            FOR (I = 0; I < N; I++)
                FOR (J = 0; J < N; J++) {
10                    INT NEWPATH = PATH[I][K] + PATH[K][J];
                    IF (NEWPATH < PATH[I][J]) PATH[I][J] =
                        NEWPATH;
                }
    }

15    Initially, PATH[I][J] is assigned EDGE[I][J], for all I
    and J. (For brevity, we use a notational shorthand for
    array assignment.) After the kth iteration, PATH[I][J] is
    the shortest path from vertex i to vertex j with
    intermediate vertices only in vertices 0 to k. Therefore,
20    after N iterations, PATH[I][J] is the shortest path from
    vertex i to vertex j with no restrictions on the
    intermediate vertices.

```

The following program solves the all-pairs shortest path problem using a multithreaded version of the Floyd-Warshall algorithm, with a barrier for thread

25 synchronization:

WO 00/36491

PCT/US99/30274

```

void ShortestPaths2(int edge[N][N], int path[N][N], int
numThreads)
{
    int t;
5    Barrier b;

    path[0..N-1][0..N-1] = edge[0..N-1][0..N-1];
    InitializeBarrier(&b, numThreads);
    multithreaded for (t = 0; t < numThreads; t++) {
10        int k, i, j;
        for (k = 0; k < N; k++) {
            for (i = t*N/numThreads; i < (t+1)*N/numThreads;
i++)
                for (j = 0; j < N; j++) {
15                    int newPath = path[i][k] + path[k][j];
                    if (newPath < path[i][j]) path[i][j] =
newPath;
                }
            PassBarrier(&b);
20        }
    }
    FinalizeBarrier(&b);
}

```

25 The multithreaded outer loop creates NUMTHREADS threads. Each thread executes the N iterations of the Floyd-Warshall algorithm on a subset of the rows of the path matrix. To keep the iterations synchronized, the threads pass through an N-way barrier at the end of each

30 iteration. There are no sharing violations on the concurrent accesses to PATH across the threads, because the algorithm will never assign to PATH[I][K] or PATH[K][J] during iteration k.

 The barrier algorithm successfully divides the work

35 among an arbitrary number of threads. However, in requiring that all threads complete each iteration before any thread begins the next iteration, the algorithm does not express the full opportunities for concurrency inherent in the data dependencies. As a consequence, the

WO 00/36491

PCT/US99/30274

program is less than optimally efficient. N-way
 synchronization at the barrier is a bottleneck that
 creates delays on entry and exit, and processor load
 imbalance can occur if all threads do not reach the
 5 barrier simultaneously.

A More Efficient Multithreaded Solution Using Condition Variable Synchronization

The following program solves the all-pairs shortest
 path problem using a more efficient multithreaded version
 10 of the Floyd-Warshall algorithm, with an array of N
 condition variables for thread synchronization:

```

void ShortestPaths3(int edge[N][N], int path[N][N], int
numThreads)
{
15   int k, t;
      Condition kDone[N];
      int kRow[N][N];

      path[0..N-1][0..N-1] = edge[0..N-1][0..N-1];
20   for (k = 0; k < N; k++) InitializeCondition(&kDone[k]);
      kRow[0] = path[0][0..N-1];
      SetCondition(&kDone[0]);
      multithreaded for (t = 0; t < numThreads; t++) {
25         int k, i, j;
         for (k = 0; k < N; k++) {
             CheckCondition(&kDone[k]);
             for (i = t*N/numThreads; i < (t+1)*N/numThreads;
i++) {
30                 for (j = 0; j < N; j++) {
                     int newPath = path[i][k] + kRow[k][j];
                     if (newPath < path[i][j]) path[i][j] =
newPath;
35                     }
                     if (i == k+1) {
                         kRow[k+1][0..N-1] = path[k+1][0..N-1];
                         SetCondition(&kDone[k+1]);
                     }
                 }
             }
40         }
      for (k = 0; k < N; k++) FinalizeCondition(&kDone[k]);
}

```


WO 00/36491

PCT/US99/30274

As with the barrier algorithm, each thread executes the N iterations of the Floyd-Warshall algorithm on a subset of the rows of the PATH matrix. However, each thread can individually continue with its next iteration
5 as soon as the necessary data is available, instead of waiting for the previous iteration to complete in all the other threads. Condition variable `KDONE[K]` is set when row k of the PATH matrix has been computed in iteration k-1. Each thread waits on `KDONE[K]` before executing iteration
10 k. To avoid sharing violations, row k of the PATH matrix computed in iteration k-1 is stored in `KRow[k]`.

The condition variable algorithm avoids the inefficiencies associated with barrier synchronization. Threads synchronize individually, rather than in an N-way
15 bottleneck, and faster threads can execute many iterations ahead of slower threads. Potentially, the N threads can be executing in up to N different iterations. One extra cost of this algorithm is the storage for the `KRow` matrix. However, the most significant extra cost is
20 allocation of N condition variables.

The following program solves the all-pairs shortest path problem using the efficient multithreaded version of the Floyd-Warshall algorithm, with a single counter for thread synchronization in place of N condition variables:

WO 00/36491

PCT/US99/30274

```

void ShortestPaths3(int edge[N][N], int path[N][N], int
numThreads)
{
    int k, t;
5    Counter kCount;
    int kRow[N][N];

    path[0..N-1][0..N-1] = edge[0..N-1][0..N-1];
    InitializeCounter(&kCount);
10    kRow[0] = path[0][0..N-1];
    multithreaded for (t = 0; t < numThreads; t++) {
        int k, i, j;
        for (k = 0; k < N; k++) {
            CheckCounter(&kCount, k);
15            for (i = t*N/numThreads; i < (t+1)*N/numThreads;
i++) {
                for (j = 0; j < N; j++) {
                    int newPath = path[i][k] + kRow[k][j];
20                    if (newPath < path[i][j]) path[i][j] =
newPath;
                }
                if (i == k+1) {
                    kRow[k+1][0..N-1] = path[k+1][0..N-1];
25                    IncrementCounter(&kCount, 1);
                }
            }
        }
        FinalizeCounter(&kCount);
30    }
}

```

Operations on N different values of the single counter replace operations on N different elements of the array of condition variables. The algorithm has the same

35 performance advantages over the barrier algorithm, without the cost of statically allocating and maintaining N synchronization objects. Internally, the counter may create synchronization objects for the distinct counter values on which threads are suspended. However, in

40 practice, the number of these objects in existence at any given time is likely to be a small fraction of N.

Three Synchronization Patterns Example

WO 00/36491

PCT/US99/30274

Three examples of practical synchronization patterns are described that can be expressed more elegantly (and often more efficiently) using counters than with traditional synchronization mechanisms. For each of these

5 synchronization patterns, a small example program is provided to demonstrate the pattern and a description of the importance of the pattern to real problems. This is far from an exhaustive list of patterns to which counters can usefully be applied. Counters are equally applicable

10 to many other situations, particularly dataflow style synchronization patterns arising in the application of threads to multiprocessing.

Counters can often be used to replace traditional barrier synchronization with a less restrictive form of

15 "ragged" barrier. With a ragged barrier, each thread waits at the barrier point only until its own individual data dependencies have been satisfied, instead of until the data dependencies of all threads have been satisfied, as with a traditional barrier. We have already given one

20 example of this pattern in Section 0, with the multithreaded Floyd-Warshall algorithm to solve the all-pairs shortest-path problem. In this section, we give another more straightforward example, based on boundary exchange in a time-stepped simulation.

WO 00/36491

PCT/US99/30274

Consider a time-stepped simulation of a one-dimensional object subdivided into N cells. The state of internal cell i at time t is a function of the states of cells $i-1$, i , and $i+1$ at time $t-1$. The states of the

5 leftmost and rightmost cells remain constant over time. An example is simulation of heat transfer along a metal rod. Similar boundary exchange requirements occur in most multithreaded simulations of physical systems in one or more dimensions. These requirements are traditionally

10 satisfied using barrier synchronization.

The following program implements the simulation using one thread for each cell, with traditional barrier synchronization between threads before cell state exchanges and updates at each time step:

```

15     float state[N];
        Barrier b;
        ...
        state[0..N-1] = initial cell states;
        InitializeBarrier(&b, N-2);
20     multithreaded for(i = 1; i < N-1; i++) {
            float leftState, rightState;
            for (t = 1; t <= numSteps; t++) {
                PassBarrier(&b);
                leftState = state[i-1];
25                rightState = state[i+1];
                PassBarrier(&b);
                state[i] = f(leftState, state[i], rightState);
            }
        }
30     FinalizeBarrier(&b);

```

All threads synchronize at the barrier twice every time step: once before exchanging cell states, and again before updating cell states. However, complete barrier

35 synchronization between all threads is unnecessarily

WO 00/36491

PCT/US99/30274

restrictive. The conditions for safely exchanging and updating the cell states involve dependencies between pairs of neighboring cells, not across all cells. As a consequence of using barriers, the performance of the

5 program is potentially subject to synchronization bottleneck and load imbalance problems.

The following program implements the same simulation using an array of counters to provide ragged barrier synchronization between threads:

```

10     float state[N];
        Counter c[N];
        ...
        state[0..N-1] = initial cell states;
        for (i = 0; i < N; i++) CounterInitialise(&c[i]);
15     IncrementCounter(&c[0], 2*numSteps);
        IncrementCounter(&c[N-1], 2*numSteps);
        multithreaded for (i = 1; i < N-1; i++) {
            float leftState, rightState, myState = state[i];
            for (t = 1; t <= numSteps; t++) {
20                 CheckCounter(&c[i-1], 2*t-2); leftState = state[i-1];
                    CheckCounter(&c[i+1], 2*t-2); rightState = state[i+1];
                    IncrementCounter(&c[i], 1);
                    myState = f(leftState, myState, rightState);
25                 CheckCounter(&c[i-1], 2*t-1);
                    CheckCounter(&c[i+1], 2*t-1);
                    state[i] = myState;
                    IncrementCounter(&c[i], 1);
            }
        }
30     for (i = 0; i < N; i++) FinalizeCounter(&c[i]);

```

As with the traditional barrier algorithm, the threads synchronize every time step before exchanging cell states, and again before updating cell states.

35 However, the synchronization is between pairs of neighboring threads via an array of counters. $c[i] = 2 \cdot t - 1$ indicates that thread i has finished reading both neighboring cell states in time step t , and $c[i] = 2 \cdot t$

WO 00/36491

PCT/US99/30274

indicates that thread i has completed time step τ .
Pairwise synchronization removes the synchronization
bottleneck of a traditional barrier and reduces load
imbalance by allowing some threads to execute ahead of
5 other threads. The barrier could be made even more ragged
using separate counters to synchronize with left and
right neighbors.

The major cost in the implementation of ragged
barriers using counters is the need for N counter objects
10 instead of one barrier object. However, the number of
counters needed is proportional to the number of threads,
not to the problem size. This cost is unlikely to be a
practical problem on modern computer systems.

The present application can be used in multithreaded
15 programming system, with any single or multiprocessor
computers. Example multithreaded programming systems
include Windows NT, UNIX/Pthreads and Java.

Other examples than those discussed above can of
course be used. While the three examples discussed above
20 are computationally intensive, other computationally
intensive systems include volume rendering, terrain
masking, threat analysis, protein folding, and molecular
dynamics simulation.

WO 00/36491

PCT/US99/30274

As can be seen from the above, the system of the present application is highly advantageous and produces significant advantages.

Although only a few embodiments have been disclosed
5 in detail above, other modifications are possible, and would understood by those having ordinary skill in the art reading the application. For example, although this application has only described certain operating systems which capable of handling multiple threads, it should be
10 understood that other operating systems could be provided. A non-exhaustive list of operating systems includes Windows NT, Windows 2000, Java, UNIX, Linux or any other type system.

All such modifications are intended to be
15 encompassed within the following claims, in which:

WO 00/36491

PCT/US99/30274

Implementation of the Sthreads library.

```

#ifndef STHREADS_H
#define STHREADS_H

#ifndef _WIN32
#error ERROR: Win32 pthreads.h included in non-Win32 program.
#endif

#ifndef _MT
#error ERROR: Sthreads program must be linked with multithreaded libraries.
#endif

#ifdef __cplusplus
extern "C" {
#endif

/*-----*/
/* Sthreads: A Structured Thread Library for Shared-Memory Multiprocessing */
/* Version 1.0 for Windows NT */
/* */
/* Author: John Thornley, Computer Science Dept., Caltech. */
/* Date: September 1998. */
/* */
/* Copyright (c) 1998 by John Thornley. */
/*-----*/

/*-----*/
/* Error codes */
/*-----*/

#define STHREADS_ERROR_NONE 0
#define STHREADS_ERROR_INPUTVALUE 1
#define STHREADS_ERROR_MEMORYALLOC 2
#define STHREADS_ERROR_THREADCREATE 3
#define STHREADS_ERROR_SYNCCREATE 4
#define STHREADS_ERROR_INITIALIZED 5
#define STHREADS_ERROR_UNINITIALIZED 6
#define STHREADS_ERROR_FINALIZED 7
#define STHREADS_ERROR_INUSE 8
#define STHREADS_ERROR_LOCKHELD 9
#define STHREADS_ERROR_LOCKNOTHELD 10
#define STHREADS_ERROR_COUNTEROVERFLOW 11
#define STHREADS_ERROR_UNSPECIFIED 12

/* Requirements: */
/* - STHREADS_ERROR_NONE == 0. */
/* - STHREADS_ERROR_INPUTVALUE > STHREADS_ERROR_NONE. */
/* - STHREADS_ERROR_MEMORYALLOC > STHREADS_ERROR_INPUTVALUE. */
/* - STHREADS_ERROR_THREADCREATE > STHREADS_ERROR_MEMORYALLOC. */
/* - STHREADS_ERROR_SYNCCREATE > STHREADS_ERROR_THREADCREATE. */
/* - STHREADS_ERROR_INITIALIZED > STHREADS_ERROR_SYNCCREATE. */
/* - STHREADS_ERROR_UNINITIALIZED > STHREADS_ERROR_INITIALIZED. */
/* - STHREADS_ERROR_FINALIZED > STHREADS_ERROR_UNINITIALIZED. */
/* - STHREADS_ERROR_INUSE > STHREADS_ERROR_FINALIZED. */
/* - STHREADS_ERROR_LOCKHELD > STHREADS_ERROR_INUSE. */
/* - STHREADS_ERROR_LOCKNOTHELD > STHREADS_ERROR_LOCKHELD. */
/* - STHREADS_ERROR_COUNTEROVERFLOW > STHREADS_ERROR_LOCKNOTHELD. */
/* - STHREADS_ERROR_UNSPECIFIED > STHREADS_ERROR_COUNTEROVERFLOW. */
/* - STHREADS_ERROR_UNSPECIFIED < INT_MAX. */

/*-----*/
/* Error string maximum length */
/*-----*/

#define STHREADS_ERROR_STRING_MAX 100

/* Requirements: */
/* - STHREADS_ERROR_STRING_MAX >= 1. */
/* - STHREADS_ERROR_STRING_MAX <= INT_MAX. */

/*-----*/
/* Processors */
/*-----*/

#define STHREADS_PROCESSORS_MAX 32

```


WO 00/36491

PCT/US99/30274

```

#define STHREADS_PROCESSOR_YES 1000
#define STHREADS_PROCESSOR_NO 1001

/* Requirements: */
/* - STHREADS_PROCESSORS_MAX >= 1. */
/* - STHREADS_PROCESSORS_MAX <= INT_MAX. */
/* - STHREADS_PROCESSOR_YES >= INT_MIN. */
/* - STHREADS_PROCESSOR_YES <= INT_MAX. */
/* - STHREADS_PROCESSOR_NO >= INT_MIN. */
/* - STHREADS_PROCESSOR_NO <= INT_MAX. */
/* - STHREADS_PROCESSOR_YES != STHREADS_PROCESSOR_NO. */

/* Definitions: */
/* - ValidProcessorStatus(p) = */
/*   p == STHREADS_PROCESSOR_PRESENT || */
/*   p == STHREADS_PROCESSOR_NOT_PRESENT. */

/*-----*/
/* Mappings of statements/iterations to threads */
/*-----*/

#define STHREADS_MAPPING_SIMPLE 3000
#define STHREADS_MAPPING_DYNAMIC 3001
#define STHREADS_MAPPING_BLOCKED 3002
#define STHREADS_MAPPING_INTERLEAVED 3003

/* Requirements: */
/* - STHREADS_MAPPING_SIMPLE > 0. */
/* - STHREADS_MAPPING_DYNAMIC == STHREADS_MAPPING_SIMPLE + 1. */
/* - STHREADS_MAPPING_BLOCKED == STHREADS_MAPPING_DYNAMIC + 1. */
/* - STHREADS_MAPPING_INTERLEAVED == STHREADS_MAPPING_BLOCKED + 1. */
/* - STHREADS_MAPPING_INTERLEAVED < INT_MAX. */

/* Definitions: */
/* - ValidMapping(m) = */
/*   m == STHREADS_MAPPING_SIMPLE || */
/*   m == STHREADS_MAPPING_DYNAMIC || */
/*   m == STHREADS_MAPPING_BLOCKED || */
/*   m == STHREADS_MAPPING_INTERLEAVED. */

/*-----*/
/* Conditions testable in regular for loop control */
/*-----*/

#define STHREADS_CONDITION_LT 4000
#define STHREADS_CONDITION_LE 4001
#define STHREADS_CONDITION_GT 4002
#define STHREADS_CONDITION_GE 4003

/* Requirements: */
/* - STHREADS_CONDITION_LT > 0. */
/* - STHREADS_CONDITION_LE == STHREADS_CONDITION_LT + 1. */
/* - STHREADS_CONDITION_GT == STHREADS_CONDITION_LE + 1. */
/* - STHREADS_CONDITION_GE == STHREADS_CONDITION_GT + 1. */
/* - STHREADS_CONDITION_GE < INT_MAX. */

/* Definitions: */
/* - ValidCondition(c) = */
/*   c == STHREADS_CONDITION_LT || */
/*   c == STHREADS_CONDITION_LE || */
/*   c == STHREADS_CONDITION_GT || */
/*   c == STHREADS_CONDITION_GE. */

/*-----*/
/* Stack sizes (in bytes) */
/*-----*/

#define STHREADS_STACK_SIZE_MINIMUM 16384
#define STHREADS_STACK_SIZE_DEFAULT 262144

/* Requirements: */
/* - STHREADS_STACK_SIZE_MINIMUM >= 0. */
/* - STHREADS_STACK_SIZE_DEFAULT >= STHREADS_STACK_SIZE_MINIMUM. */
/* - STHREADS_STACK_SIZE_DEFAULT <= UINT_MAX. */

```

WO 00/36491

PCT/US99/30274

```

/* Definitions: */
/* - ValidStackSize(s) = */
/*   s >= STHREADS_STACK_SIZE_MINIMUM. */

/*-----*/
/* Priorities */
/*-----*/

#define STHREADS_PRIORITY_LOWEST -2
#define STHREADS_PRIORITY_HIGHEST +2
#define STHREADS_PRIORITY_PARENT 10000 /* Inherit priority of parent thread. */

/* Requirements: */
/* - STHREADS_PRIORITY_LOWEST > INT_MIN. */
/* - STHREADS_PRIORITY_HIGHEST >= STHREADS_PRIORITY_LOWEST. */
/* - STHREADS_PRIORITY_HIGHEST < INT_MAX. */
/* - STHREADS_PRIORITY_PARENT < STHREADS_PRIORITY_LOWEST || */
/*   STHREADS_PRIORITY_PARENT > STHREADS_PRIORITY_HIGHEST. */

/* Definitions: */
/* - ValidPriority(p) = */
/*   STHREADS_PRIORITY_LOWEST <= p && p <= STHREADS_PRIORITY_HIGHEST. */

/*-----*/
/* Print error message to string */
/*-----*/

void SthreadsWriteErrorMessage(int errorCode, char errorString[]);

/* Input Arguments: */
/* - errorCode : error code returned by an Sthreads function call. */
/* Output Arguments: */
/* - errorString : error message as a char string. */
/* Preconditions: */
/* - errorString != NULL && */
/*   errorString is a string of at least STHREADS_ERROR_STRING_MAX chars. */
/* Postconditions: */
/* - errorString is '\0' terminated string of chars in the range ' ' .. '-'. */
/* - 1 <= strlen(errorString) < STHREADS_ERROR_STRING_MAX. */
/* Atomicity: */
/* - Atomic with respect to all operations. */

/*-----*/
/* Handle errors. */
/*-----*/

void SthreadsErrorHandler(int errorCode);

/* Input Arguments: */
/* - errorCode : error code returned by an Sthreads function call. */
/* Operation: */
/* - error handler function is called with errorCode as argument. */
/* Default Error Handler Function: */
/* - Displays error message and terminates normal program execution. */
/* Atomicity: */
/* - Not atomic with respect to SthreadsSetErrorHandler operations. */
/* - Atomic with respect to all other operations. */

/*-----*/
/* Set error handler function. */
/*-----*/

int SthreadsSetErrorHandler(void (*errorHandler)(int errorCode));

/* Input Arguments: */
/* - errorHandler : function to handle errors. */
/* Preconditions: */
/* - errorHandler == NULL || */
/*   errorHandler is valid void (*)(int) function. */
/* Postconditions: */
/* - if (errorHandler == NULL) */
/*   error handler function is set to default error handler function. */
/* - if (errorHandler != NULL) */

```

WO 00/36491

PCT/US99/30274

```

/*      error handler function is set to ErrorHandler.          */
/* Atomicity:                                                    */
/* - Not atomic with respect to                                  */
/*   SthreadsHandleError and SthreadsSetErrorHandler operations. */
/* - Atomic with respect to all other operations.                */

/*-----*/
/* Control the processors used by program execution.             */
/*-----*/

int SthreadsGetSystemProcessors(int processor[]);

/* Output Arguments:                                             */
/* - processors : processors that exist on the system.           */
/* Function Return:                                              */
/* - error code.                                                 */
/* Preconditions:                                                */
/* - processor != NULL &&                                         */
/*   processor is an array of at least STHREADS_PROCESSORS_MAX ints. */
/* Postconditions:                                               */
/* - forall (p = 0; p < STHREADS_PROCESSORS_MAX; p++)           */
/*   ValidProcessorStatus(processor[p]) &&                       */
/*   (if (processor[p] == STHREADS_PROCESSOR_YES)                */
/*     a processor numbered p exists on the system) &&         */
/*   (if (processor[p] == STHREADS_PROCESSOR_NO)                 */
/*     a processor numbered p does not exist on the system).    */
/* Atomicity:                                                    */
/* - Atomic with respect to all operations.                      */

int SthreadsSetProgramProcessors(int processor[]);

/* Input Arguments:                                             */
/* - processor : processors on which the threads of the program may execute. */
/* Function Return:                                              */
/* - error code.                                                 */
/* Preconditions:                                                */
/* - processor != NULL &&                                         */
/*   processor is an array of at least STHREADS_PROCESSORS_MAX ints. */
/* - forall (p = 0; p < STHREADS_PROCESSORS_MAX; p++)           */
/*   ValidProcessorStatus(processor[p]) &&                       */
/*   if (processor[p] == STHREADS_PROCESSOR_YES)                */
/*     a processor numbered p exists on the system.            */
/* - exists (p = 0; p < STHREADS_PROCESSORS_MAX; p++)           */
/*   processor[p] == STHREADS_PROCESSOR_YES.                     */
/* Atomicity:                                                    */
/* - Must be called when program execution consists of a single thread. */

int SthreadsGetProgramProcessors(int processor[]);

/* Output Arguments:                                             */
/* - processors : processors on which the program may execute.   */
/* Function Return:                                              */
/* - error code.                                                 */
/* Preconditions:                                                */
/* - processor != NULL &&                                         */
/*   processor is an array of at least STHREADS_PROCESSORS_MAX ints. */
/* Postconditions:                                               */
/* - forall (p = 0; p < STHREADS_PROCESSORS_MAX; p++)           */
/*   ValidProcessorStatus(processor[p]) &&                       */
/*   (if (processor[p] == STHREADS_PROCESSOR_YES)                */
/*     the program may execute on processor number p) &&         */
/*   (if (processor[p] == STHREADS_PROCESSOR_NO)                 */
/*     the program may not execute on processor number p).      */
/* Atomicity:                                                    */
/* - Not atomic with respect to                                  */
/*   SetProgramProcessors and SetNumProgramProcessors operations. */
/* - Atomic with respect to all other operations.                */

int SthreadsSetThreadProcessors(int processor[]);

/* Input Arguments:                                             */
/* - processor : processors on which the thread may execute.     */
/* Function Return:                                              */
/* - error code.                                                 */

```

WO 00/36491

PCT/US99/30274

```

/* Preconditions:
/* - processor != NULL &&
/* - processor is an array of at least STHREADS_PROCESSORS_MAX ints.
/* - forall (p = 0; p < STHREADS_PROCESSORS_MAX; p++)
/*     ValidProcessorStatus(processor[p]) &&
/*     if (processor[p] == STHREADS_PROCESSOR_YES)
/*         the program may execute on processor number p.
/* - exists (p = 0; p < STHREADS_PROCESSORS_MAX; p++)
/*     processor[p] == STHREADS_PROCESSOR_YES.
/* Atomicity:
/* - Not atomic with respect to
/* - SetProgramProcessors and SetNumProgramProcessors operations.
/* - Atomic with respect to all other operations.

int SthreadsGetNumSystemProcessors(int *numProcessors);

/* Output Arguments:
/* - numProcessors : number of processors that exist on the system.
/* Function Return:
/* - error code.
/* Preconditions:
/* - numProcessors != NULL && numProcessors points to a valid int variable.
/* Postconditions:
/* - *numProcessors == number of processors that exist on the system.
/* Atomicity:
/* - Atomic with respect to all operations.

int SthreadsSetNumProgramProcessors(int numProcessors);

/* Input Arguments:
/* - numProcessors : number of processors on which the threads of the program
/*                   may execute.
/* Function Return:
/* - error code.
/* Preconditions:
/* - numProcessors >= 1.
/* - numProcessors <= number of processors that exist on the system.
/* Atomicity:
/* - Must be called when program execution consists of a single thread.

/*-----*/
/* Multithreaded block
/*-----*/

int SthreadsBlock(
    int numStatements, void (*statement[]) (void *args), void *args,
    int mapping, int numThreads,
    int priority, unsigned int stackSize);

/* Input Arguments:
/* - numStatements : number of statements in block.
/* - statement     : functions representing statements.
/* - args          : pointer to arguments of the statements.
/* - mapping       : mapping of statements onto threads.
/* - numThreads    : number of threads.
/* - priority      : priority of threads.
/* - stackSize     : stack size of threads.
/* Function Return:
/* - error code.
/* Preconditions:
/* - numStatements >= 0.
/* - statement != NULL &&
/*   statement is an array of at least numStatements functions.
/* - forall (s = 0; s < numStatements; s++)
/*     statement[s] != NULL &&
/*     statement[s] is a valid void (*) (void *) function.
/* - ValidMapping(mapping).
/* - if (mapping != STHREADS_MAPPING_SIMPLE)
/*     (numThreads > 0) || (numThreads == 0 && numStatements == 0).
/* - ValidPriority(priority) || priority == STHREADS_PRIORITY_PARENT.
/* - ValidStackSize(stackSize).
/* Atomicity:
/* - Atomic with respect to all operations.

```

WO 00/36491

PCT/US99/30274

```

/*-----*/
/* Multithreaded regular for loop */
/*-----*/

int SthreadsRegularForLoop(
    void (*chunk)(int initial, int bound, int step, void *args), void *args,
    int initial, int condition, int bound, int step,
    int chunkSize, int mapping, int numThreads,
    int priority, unsigned int stackSize);

/* Input Arguments: */
/* - chunk      : function to execute iterations of loop body. */
/* - args       : pointer to arguments of loop body. */
/* - initial    : initial value of control variable. */
/* - condition  : condition between control variable and bound value. */
/* - bound      : bound value of control variable. */
/* - step       : step value of control variable. */
/* - chunkSize  : number of iterations per chunk. */
/* - mapping    : mapping of chunks onto threads. */
/* - numThreads : number of threads. */
/* - priority   : priority of threads. */
/* - stackSize  : stack size of threads. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - chunk != NULL && */
/*   chunk is a valid void (*)(int, int, int, void *) function. */
/* - ValidCondition(condition). */
/* - !InfiniteRange(initial, condition, bound, step). */
/* - (chunkSize > 0) || */
/*   (chunkSize == 0 && NullRange(initial, condition, bound, step)). */
/* - ValidMapping(mapping). */
/* - if (mapping != STHREADS_MAPPING_SIMPLE) */
/*   (numThreads > 0) || */
/*   (numThreads == 0 && NullRange(initial, condition, bound, step)). */
/* - ValidPriority(priority) || priority == STHREADS_PRIORITY_PARENT. */
/* - ValidStackSize(stackSize). */

/* Definitions: */
/* - InfiniteRange(initial, condition, bound, step) = */
/*   (condition == STHREADS_CONDITION_LT && */
/*     initial < bound && step <= 0) || */
/*   (condition == STHREADS_CONDITION_LE && */
/*     initial <= bound && step <= 0) || */
/*   (condition == STHREADS_CONDITION_GT && */
/*     initial > bound && step >= 0) || */
/*   (condition == STHREADS_CONDITION_GE && */
/*     initial >= bound && step >= 0). */
/* - NullRange(initial, condition, bound, step) = */
/*   (condition == STHREADS_CONDITION_LT && initial >= bound) || */
/*   (condition == STHREADS_CONDITION_LE && initial > bound) || */
/*   (condition == STHREADS_CONDITION_GT && initial <= bound) || */
/*   (condition == STHREADS_CONDITION_GE && initial < bound). */
/* Atomicity: */
/* - Atomic with respect to all operations. */

/*-----*/
/* Flags */
/*-----*/

typedef struct {
    unsigned char value[16];
} SthreadsFlag;

int SthreadsFlagInitialize(SthreadsFlag *flag);

/* Input-Output Arguments: */
/* - flag : flag variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - flag != NULL && flag points to a valid flag variable. */
/* - !Initialized(flag). */
/* Atomicity: */

```

WO 00/36491

PCT/US99/30274

```

/* - Not atomic with respect to all other operations on flag.
/* - Atomic with respect to all other operations.
*/

int SthreadsFlagFinalize(SthreadsFlag *flag);

/* Input-Output Arguments:
/* - flag : flag variable.
/* Function Return:
/* - error code.
/* Preconditions:
/* - flag != NULL && flag points to a valid flag variable.
/* - Initialized(flag) && !Finalized(flag).
/* - NumWaiting(flag) == 0.
/* Atomicity:
/* - Not atomic with respect to all other operations on flag.
/* - Atomic with respect to all other operations.
*/

int SthreadsFlagSet(SthreadsFlag *flag);

/* Input-Output Arguments:
/* - flag : flag variable.
/* Function Return:
/* - error code.
/* Preconditions:
/* - flag != NULL && flag points to a valid flag variable.
/* - Initialized(flag) && !Finalized(flag).
/* Atomicity:
/* - Atomic with respect to Set and Check operations on flag.
/* - Not atomic with respect to other operations on flag.
/* - Atomic with respect to all other operations.
*/

int SthreadsFlagCheck(SthreadsFlag *flag);

/* Input-Output Arguments:
/* - flag : flag variable.
/* Function Return:
/* - error code.
/* Preconditions:
/* - flag != NULL && flag points to a valid flag variable.
/* - Initialized(flag) && !Finalized(flag).
/* Atomicity:
/* - Atomic with respect to Set and Check operations on flag.
/* - Not atomic with respect to other operations on flag.
/* - Atomic with respect to all other operations.
*/

int SthreadsFlagReset(SthreadsFlag *flag);

/* Input-Output Arguments:
/* - flag : flag variable.
/* Function Return:
/* - error code.
/* Preconditions:
/* - flag != NULL && flag points to a valid flag variable.
/* - Initialized(flag) && !Finalized(flag).
/* - NumWaiting(flag) == 0.
/* Atomicity:
/* - Not atomic with respect to other operations on flag.
/* - Atomic with respect to all other operations.
*/

/*-----*/
/* Counters
/*-----*/

typedef struct {
    unsigned char value[40];
} SthreadsCounter;

int SthreadsCounterInitialize(SthreadsCounter *counter);

/* Input-Output Arguments:
/* - counter : pointer to counter variable.
/* Function Return:
/* - error code.
/* Preconditions:
*/

```

WO 00/36491

PCT/US99/30274

```

/* - counter != NULL && counter points to a valid counter variable. */
/* - !Initialized(counter). */
/* Atomicity: */
/* - Not atomic with respect to all other operations on counter. */
/* - Atomic with respect to all other operations. */

int SthreadsCounterFinalize(SthreadsCounter *counter);

/* Input-Output Arguments: */
/* - counter : pointer to counter variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - counter != NULL && counter points to a valid counter variable. */
/* - Initialized(counter) && !Finalized(counter). */
/* - NumWaiting(counter) == 0. */
/* Atomicity: */
/* - Not atomic with respect to all other operations on counter. */
/* - Atomic with respect to all other operations. */

int SthreadsCounterIncrement(SthreadsCounter *counter, unsigned int amount);

/* Input-Output Arguments: */
/* - counter : pointer to counter variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - counter != NULL && counter points to a valid counter variable. */
/* - Initialized(counter) && !Finalized(counter). */
/* - Count(counter) <= UINT_MAX - amount. */
/* Atomicity: */
/* - Atomic with respect to Increment and Check operations on counter. */
/* - Not atomic with respect to other operations on counter. */
/* - Atomic with respect to all other operations. */

int SthreadsCounterCheck(SthreadsCounter *counter, unsigned int value);

/* Input-Output Arguments: */
/* - counter : pointer to counter variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - counter != NULL && counter points to a valid counter variable. */
/* - Initialized(counter) && !Finalized(counter). */
/* Atomicity: */
/* - Atomic with respect to Increment and Check operations on counter. */
/* - Not atomic with respect to other operations on counter. */
/* - Atomic with respect to all other operations. */

int SthreadsCounterReset(SthreadsCounter *counter);

/* Input-Output Arguments: */
/* - counter : pointer to counter variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - counter != NULL && counter points to a valid counter variable. */
/* - Initialized(counter) && !Finalized(counter). */
/* - NumWaiting(counter) == 0. */
/* Atomicity: */
/* - Not atomic with respect to all other operations on counter. */
/* - Atomic with respect to all other operations. */

/*-----*/
/* Locks */
/*-----*/

typedef struct {
    unsigned char value[36];
} SthreadsLock;

int SthreadsLockInitialize(SthreadsLock *lock);

/* Input-Output Arguments: */

```

WO 00/36491

PCT/US99/30274

```

/* - lock : pointer to lock variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - lock != NULL && lock points to a valid lock variable. */
/* - !Initialized(lock). */
/* Atomicity: */
/* - Not atomic with respect to all other operations on lock. */
/* - Atomic with respect to all other operations. */

int SthreadsLockFinalize(SthreadsLock *lock);

/* Input-Output Arguments: */
/* - lock : pointer to lock variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - lock != NULL && lock points to a valid lock variable. */
/* - Initialized(lock) && !Finalized(lock). */
/* - !AnyThreadHolds(lock). */
/* Atomicity: */
/* - Not atomic with respect to all other operations on lock. */
/* - Atomic with respect to all other operations. */

int SthreadsLockAcquire(SthreadsLock *lock);

/* Input-Output Arguments: */
/* - lock : pointer to lock variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - lock != NULL && lock points to a valid lock variable. */
/* - Initialized(lock) && !Finalized(lock). */
/* - !ThisThreadHolds(lock). */
/* Atomicity: */
/* - Atomic with respect to Acquire and Release operations on lock. */
/* - Not atomic with respect to other operations on lock. */
/* - Atomic with respect to all other operations. */

int SthreadsLockRelease(SthreadsLock *lock);

/* Input-Output Arguments: */
/* - lock : pointer to lock variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - lock != NULL && lock points to a valid lock variable. */
/* - Initialized(lock) && !Finalized(lock). */
/* - ThisThreadHolds(lock). */
/* Atomicity: */
/* - Atomic with respect to Acquire and Release operations on lock. */
/* - Not atomic with respect to other operations on lock. */
/* - Atomic with respect to all other operations. */

/*-----*/
/* Barriers */
/*-----*/

typedef struct {
    unsigned char value[52];
} SthreadsBarrier;

int SthreadsBarrierInitialize(SthreadsBarrier *barrier, int numThreads);

/* Input-Output Arguments: */
/* - barrier : pointer to barrier variable. */
/* - numThreads : number of threads that cross barrier in each pass. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - barrier != NULL && barrier points to a valid barrier variable. */
/* - !Initialized(barrier). */
/* - numThreads >= 1. */
/* Atomicity: */

```


WO 00/36491

PCT/US99/30274

```

/* - Not atomic with respect to all other operations on barrier. */
/* - Atomic with respect to all other operations. */

int SthreadsBarrierFinalize(SthreadsBarrier *barrier);

/* Input-Output Arguments: */
/* - barrier : pointer to barrier variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - barrier != NULL && barrier points to a valid barrier variable. */
/* - Initialized(barrier) && !Finalized(barrier). */
/* - NumWaiting(barrier) == 0. */
/* Atomicity: */
/* - Not atomic with respect to all other operations on barrier. */
/* - Atomic with respect to all other operations. */

int SthreadsBarrierPass(SthreadsBarrier *barrier);

/* Input-Output Arguments: */
/* - barrier : pointer to barrier variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - barrier != NULL && barrier points to a valid barrier variable. */
/* - Initialized(barrier) && !Finalized(barrier). */
/* Atomicity: */
/* - Atomic with respect to Pass operations on barrier. */
/* - Not atomic with respect to other operations on barrier. */
/* - Atomic with respect to all other operations. */

int SthreadsBarrierReset(SthreadsBarrier *barrier, int numThreads);

/* Input-Output Arguments: */
/* - barrier : pointer to barrier variable. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - barrier != NULL && barrier points to a valid barrier variable. */
/* - Initialized(barrier) && !Finalized(barrier). */
/* - NumWaiting(barrier) == 0. */
/* - numThreads >= 1. */
/* Atomicity: */
/* - Not atomic with respect to all other operations on barrier. */
/* - Atomic with respect to all other operations. */

/*-----*/
/* Priorities */
/*-----*/

int SthreadsGetCurrentPriority(int *priority);

/* Output Arguments: */
/* - priority : scheduling priority of calling thread. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - priority != NULL && priority points to a valid int variable. */
/* Postconditions: */
/* - *priority == scheduling priority of calling thread. */
/* Atomicity: */
/* - Atomic with respect to all operations. */

int SthreadsSetCurrentPriority(int priority);

/* Input Arguments: */
/* - priority : scheduling priority for calling thread. */
/* Function Return: */
/* - error code. */
/* Preconditions: */
/* - ValidPriority(priority). */
/* Atomicity: */
/* - Atomic with respect to all operations. */

```

WO 00/36491

PCT/US99/30274

```
/*-----  
#ifdef __cplusplus  
}  
#endif  
#endif /* !STHEADS_H */
```

WO 00/36491

PCT/US99/30274

```

/*-----*/
/* Sthreads: A Structured Thread Library for Shared-Memory Multiprocessing */
/* Version 1.0 for Windows NT */
/*
/* Author: John Thornley, Computer Science Dept., Caltech.
/* Date: September 1998.
/*
/* Copyright (c) 1998 by John Thornley.
/*
/*
/* THINGS TO DO:
/*
/* - Change names of CHECK tests, e.g., to CHECKNOTINITIALIZED.
/* - Make Finalize operations set Initialized and Finalized flags to false.
/* - Counter for dynamic for loop should be unsigned int.
/* - Declarations of thread functions should be compatible with
/* Win32 prototype ... see page 25.
/* - Implement special case of BarrierPass when numThreads == 1.
/* - Implement flags like counters for efficiency when flag is set?
/* - Change priority low and high to THREAD_PRIORITY_IDLE and _TIME_CRITICAL.
/*
/*-----*/

#include <stddef.h>
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <limits.h>
#include <windows.h>
#include "sthreads.h"

/*-----*/
/* Bool type definition */
/*-----*/

typedef int bool;
#define false 0
#define true 1

/*-----*/
/* Miscellaneous utility definitions */
/*-----*/

#define MIN(x, y) ((x) < (y) ? (x) : (y))
#define MAX(x, y) ((x) > (y) ? (x) : (y))

/*-----*/
/* Verify requirements, beliefs, and checks */
/*-----*/

#define require(condition) assert(condition) /* require this input condition */
#define believe(condition) assert(condition) /* believe this must be true */
#define check(condition) assert(condition) /* check this is true */

/*-----*/
/* Check for error conditions */
/*-----*/

#define CHECKINPUTVALUE(condition) \
    if (!(condition)) { return STHREADS_ERROR_INPUTVALUE; }

#define CHECKMEMORYALLOC(condition) \
    if (!(condition)) { return STHREADS_ERROR_MEMORYALLOC; }

#define CHECKTHREADCREATE(condition) \
    if (!(condition)) { return STHREADS_ERROR_THREADCREATE; }

#define CHECKSYNCCREATE(condition) \
    if (!(condition)) { return STHREADS_ERROR_SYNCCREATE; }

#define CHECKINITIALIZED(condition) \
    if (!(condition)) { return STHREADS_ERROR_INITIALIZED; }

#define CHECKUNINITIALIZED(condition) \
    if (!(condition)) { return STHREADS_ERROR_UNINITIALIZED; }

```

WO 00/36491

PCT/US99/30274

```

#define CHECKFINALIZED(condition) \
    if (!(condition)) { return STHREADS_ERROR_FINALIZED; }

#define CHECKINUSE(condition) \
    if (!(condition)) { return STHREADS_ERROR_INUSE; }

#define CHECKLOCKHELD(condition) \
    if (!(condition)) { return STHREADS_ERROR_LOCKHELD; }

#define CHECKLOCKNOTHELD(condition) \
    if (!(condition)) { return STHREADS_ERROR_LOCKNOTHELD; }

#define CHECKCOUNTEROVERFLOW(condition) \
    if (!(condition)) { return STHREADS_ERROR_COUNTEROVERFLOW; }

#define CHECKOTHER(condition) \
    if (!(condition)) { return STHREADS_ERROR_UNSPECIFIED; }

/*-----*/
/* Is processor status value valid? */
/*-----*/

static bool ValidProcessorStatus(int p)
{
    return
        p == STHREADS_PROCESSOR_YES ||
        p == STHREADS_PROCESSOR_NO;
}

/*-----*/
/* Is mapping value valid? */
/*-----*/

static bool ValidMapping(int m)
{
    return
        m == STHREADS_MAPPING_SIMPLE ||
        m == STHREADS_MAPPING_DYNAMIC ||
        m == STHREADS_MAPPING_BLOCKED ||
        m == STHREADS_MAPPING_INTERLEAVED;
}

/*-----*/
/* Is condition value valid? */
/*-----*/

static bool ValidCondition(int c)
{
    return
        c == STHREADS_CONDITION_LT ||
        c == STHREADS_CONDITION_LE ||
        c == STHREADS_CONDITION_GT ||
        c == STHREADS_CONDITION_GE;
}

/*-----*/
/* Is stack-size value valid? */
/*-----*/

static bool ValidStackSize(unsigned int s)
{
    return
        s >= STHREADS_STACK_SIZE_MINIMUM;
}

/*-----*/
/* Is priority value valid? */
/*-----*/

static bool ValidPriority(int p)
{
    return
        STHREADS_PRIORITY_LOWEST <= p && p <= STHREADS_PRIORITY_HIGHEST;
}

```

WO 00/36491

PCT/US99/30274

```

,
/*-----*/
/* Print error message to string */
/*-----*/

void SthreadsWriteErrorMessage(int errorCode, char errorString[])
{
    switch (errorCode) {
        case STHREADS_ERROR_NONE:
            sprintf(errorString,
                "no error");
            break;
        case STHREADS_ERROR_INPUTVALUE:
            sprintf(errorString,
                "input value precondition violation");
            break;
        case STHREADS_ERROR_MEMORYALLOC:
            sprintf(errorString,
                "memory allocation failure");
            break;
        case STHREADS_ERROR_THREADCREATE:
            sprintf(errorString,
                "system thread creation failure");
            break;
        case STHREADS_ERROR_SYNCCREATE:
            sprintf(errorString,
                "system synchronization creation failure");
            break;
        case STHREADS_ERROR_INITIALIZED:
            sprintf(errorString,
                "initialization on previously initialized object");
            break;
        case STHREADS_ERROR_UNINITIALIZED:
            sprintf(errorString,
                "operation on uninitialized object");
            break;
        case STHREADS_ERROR_FINALIZED:
            sprintf(errorString,
                "operation on finalized object");
            break;
        case STHREADS_ERROR_INUSE:
            sprintf(errorString,
                "finalization/reset on in-use object");
            break;
        case STHREADS_ERROR_LOCKNOTHELD:
            sprintf(errorString,
                "release on lock not held");
            break;
        case STHREADS_ERROR_COUNTEROVERFLOW:
            sprintf(errorString,
                "counter overflow");
            break;
        case STHREADS_ERROR_UNSPECIFIED:
            sprintf(errorString,
                "unspecified error");
            break;
        default:
            sprintf(errorString,
                ">>>> unknown error code <<<<");
            break;
    }
}

/*-----*/
/* Default error handler function: */
/* displays error message and terminate normal program execution. */
/*-----*/

static void DefaultErrorHandler(int errorCode)
{
    char errorString[STHREADS_ERROR_STRING_MAX];

    if (errorCode != STHREADS_ERROR_NONE) {

```

WO 00/36491

PCT/US99/30274

```

        SthreadsWriteErrorMessage(errorCode, errorString);
        fprintf(stderr, "\n%s\n", errorString);
        exit(EXIT_FAILURE);
    }
}

/*-----*/
/* Error handler function. */
/*-----*/

static void (*errorHandlerFunction)(int errorCode) = DefaultErrorHandler;

/*-----*/
/* Handle errors. */
/*-----*/

#define UNLOCKED 0
#define LOCKED 1

static LONG lock = UNLOCKED;

void SthreadsErrorHandler(int errorCode)
{
    while (InterlockedExchange((LPLONG) &lock, LOCKED) != UNLOCKED);
    (*errorHandlerFunction)(errorCode);
    InterlockedExchange((LPLONG) &lock, UNLOCKED);
}

#undef UNLOCKED
#undef LOCKED

/*-----*/
/* Set error handler function. */
/*-----*/

int SthreadsSetErrorHandler(void (*errorHandler)(int errorCode))
{
    if (errorHandler == NULL)
        errorHandlerFunction = DefaultErrorHandler;
    else
        errorHandlerFunction = errorHandler;

    return STHREADS_ERROR_NONE;
}

/*-----*/
/* Control the processors used by program execution. */
/*-----*/

int SthreadsGetSystemProcessors(int processor[])
{
    DWORD processAffinity, systemAffinity, processorBit;
    int p;

    require(STHEADS_PROCESSORS_MAX == 32);
    GetProcessAffinityMask(
        GetCurrentProcess(),
        (LPDWORD) &processAffinity, (LPDWORD) &systemAffinity);

    CHECKINPUTVALUE(processor != NULL);

    processorBit = (DWORD) 1;
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
        if (systemAffinity & processorBit)
            processor[p] = STHREADS_PROCESSOR_YES;
        else
            processor[p] = STHREADS_PROCESSOR_NO;
        processorBit = processorBit << 1;
    }

    return STHREADS_ERROR_NONE;
}

/*-----*/

```

WO 00/36491

PCT/US99/30274

```

int SthreadsSetProgramProcessors(int processor[])
{
    DWORD processAffinity, systemAffinity, processorBit;
    int p;

    require(STHREADS_PROCESSORS_MAX == 32);
    GetProcessAffinityMask(
        GetCurrentProcess(),
        (LPDWORD) &processAffinity, (LPDWORD) &systemAffinity);

    CHECKINPUTVALUE(processor != NULL);
    processorBit = (DWORD) 1;
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
        CHECKINPUTVALUE(ValidProcessorStatus(processor[p]));
        if (processor[p] == STHREADS_PROCESSOR_YES)
            CHECKINPUTVALUE(systemAffinity & processorBit);
        processorBit = processorBit << 1;
    }
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++)
        if (processor[p] == STHREADS_PROCESSOR_YES) break;
    CHECKINPUTVALUE(p < STHREADS_PROCESSORS_MAX);

    processAffinity = (DWORD) 0;
    processorBit = (DWORD) 1;
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
        if (processor[p] == STHREADS_PROCESSOR_YES)
            processAffinity = processAffinity | processorBit;
        processorBit = processorBit << 1;
    }
    SetProcessAffinityMask(GetCurrentProcess(), processAffinity);
    SetThreadAffinityMask(GetCurrentThread(), processAffinity);

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsGetProgramProcessors(int processor[])
{
    DWORD processAffinity, systemAffinity, processorBit;
    int p;

    require(STHREADS_PROCESSORS_MAX == 32);
    GetProcessAffinityMask(
        GetCurrentProcess(),
        (LPDWORD) &processAffinity, (LPDWORD) &systemAffinity);

    CHECKINPUTVALUE(processor != NULL);

    processorBit = (DWORD) 1;
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
        if (processAffinity & processorBit)
            processor[p] = STHREADS_PROCESSOR_YES;
        else
            processor[p] = STHREADS_PROCESSOR_NO;
        processorBit = processorBit << 1;
    }

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsSetThreadProcessors(int processor[])
{
    DWORD threadAffinity, processAffinity, systemAffinity, processorBit;
    int p;

    require(STHREADS_PROCESSORS_MAX == 32);
    GetProcessAffinityMask(
        GetCurrentProcess(),
        (LPDWORD) &processAffinity, (LPDWORD) &systemAffinity);

```

WO 00/36491

PCT/US99/30274

```

CHECKINPUTVALUE(processor != NULL);
processorBit = (DWORD) 1;
for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
    CHECKINPUTVALUE(ValidProcessorStatus(processor[p]));
    if (processor[p] == STHREADS_PROCESSOR_YES)
        CHECKINPUTVALUE(processAffinity & processorBit);
    processorBit = processorBit << 1;
}
for (p = 0; p < STHREADS_PROCESSORS_MAX; p++)
    if (processor[p] == STHREADS_PROCESSOR_YES) break;
CHECKINPUTVALUE(p < STHREADS_PROCESSORS_MAX);

threadAffinity = (DWORD) 0;
processorBit = (DWORD) 1;
for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
    if (processor[p] == STHREADS_PROCESSOR_YES)
        threadAffinity = threadAffinity | processorBit;
    processorBit = processorBit << 1;
}
SetThreadAffinityMask(GetCurrentThread(), threadAffinity);

return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsGetNumSystemProcessors(int *numProcessors)
{
    DWORD processAffinity, systemAffinity, processorBit;
    int p, count;

    require(STHREADS_PROCESSORS_MAX == 32);
    GetProcessAffinityMask(
        GetCurrentProcess(),
        (LPDWORD) &processAffinity, (LPDWORD) &systemAffinity);

    CHECKINPUTVALUE(numProcessors != NULL);

    count = 0;
    processorBit = (DWORD) 1;
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
        if (systemAffinity & processorBit)
            count = count + 1;
        processorBit = processorBit << 1;
    }
    *numProcessors = count;

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsSetNumProgramProcessors(int numProcessors)
{
    DWORD processAffinity, systemAffinity, processorBit;
    int p, numSystemProcessors;

    require(STHREADS_PROCESSORS_MAX == 32);
    GetProcessAffinityMask(
        GetCurrentProcess(),
        (LPDWORD) &processAffinity, (LPDWORD) &systemAffinity);

    CHECKINPUTVALUE(numProcessors >= 1);
    numSystemProcessors = 0;
    processorBit = (DWORD) 1;
    for (p = 0; p < STHREADS_PROCESSORS_MAX; p++) {
        if (systemAffinity & processorBit)
            numSystemProcessors = numSystemProcessors + 1;
        processorBit = processorBit << 1;
    }
    CHECKINPUTVALUE(numProcessors <= numSystemProcessors);

    processAffinity = (DWORD) 0;
    processorBit = (DWORD) 1;

```


WO 00/36491

PCT/US99/30274

```

    for (p = 0; p < STHREADS_PROCESSORS_MAX && numProcessors > 0; p++) {
        if (systemAffinity & processorBit) {
            processAffinity = processAffinity | processorBit;
            numProcessors = numProcessors - 1;
        }
        processorBit = processorBit << 1;
    }
    believe(numProcessors == 0);
    SetProcessAffinityMask(GetCurrentProcess(), processAffinity);

    return STHREADS_ERROR_NONE;
}

/*-----*/
/* Arguments for multithreaded block thread */
/*-----*/

typedef struct {
    int numStatements;
    void (**statement)(void *args);
    void *args;
    int first, last, step;
    int *counter;
    LPCRITICAL_SECTION counterLock;
    LPLONG threadCount;
    HANDLE threadsFinished;
} MTBargs;

/*-----*/
/* Simple multithreaded block thread */
/*-----*/

static void SMTBthread(MTBargs *args)
{
    BOOL returnOK;

    require(args != NULL);
    require(args->numStatements > 0);
    require(args->statement != NULL);
    require(0 <= args->first && args->first < args->numStatements);
    require(*args->statement[args->first] != NULL);

    (*args->statement[args->first])(args->args);

    if (InterlockedDecrement(&args->threadCount) == 0) {
        returnOK = SetEvent(args->threadsFinished);
        check(returnOK);
    }
}

/*-----*/
/* Dynamic multithreaded block thread */
/*-----*/

static void DMTBthread(MTBargs *args)
{
    int s;
    bool finished;
    BOOL returnOK;

    require(args != NULL);
    require(args->numStatements > 0);
    require(args->statement != NULL);
    require(0 <= args->first && args->first < args->numStatements);
    require(args->counter != NULL);
    require(args->counterLock != NULL);

    s = args->first;
    while (true) {
        require(args->statement[s] != NULL);
        (*args->statement[s])(args->args);
        EnterCriticalSection(&args->counterLock);
        finished = (*args->counter == args->numStatements - 1);
        if (!finished) {

```

WO 00/36491

PCT/US99/30274

```

        *args->counter = *args->counter + 1;
        s = *args->counter;
    }
    LeaveCriticalSection(&args->counterLock);
    if (finished) break;
}

if (InterlockedDecrement(&args->threadCount) == 0) {
    returnOK = SetEvent(&args->threadsFinished);
    check(returnOK);
}
}

/*-----*/
/* Blocked and interleaved multithreaded block thread */
/*-----*/

static void BMTBthread(MTBargs *args)
{
    int s;
    BOOL returnOK;

    require(args != NULL);
    require(args->numStatements > 0);
    require(args->statement != NULL);
    require(0 <= args->last && args->last < args->numStatements);
    require(0 <= args->first && args->first <= args->last);
    require(args->step > 0);
    require((args->last - args->first) % args->step == 0);

    s = args->first;
    while (true) {
        require(args->statement[s] != NULL);
        (*args->statement[s])(args->args);
        if (s == args->last) break;
        believe(args->last - s >= args->step);
        s = s + args->step;
    }

    if (InterlockedDecrement(&args->threadCount) == 0) {
        returnOK = SetEvent(&args->threadsFinished);
        check(returnOK);
    }
}

/*-----*/
/* Multithreaded block */
/*-----*/

int SthreadsBlock(
    int numStatements, void (*statement[])(void *args), void *args,
    int mapping, int numThreads,
    int priority, unsigned int stackSize)
{
    HANDLE *thread;
    MTBargs *threadArgs;
    LONG threadCount;
    HANDLE threadsFinished;
    HANDLE parentThread;
    int parentPriority;
    void (*threadStart)(MTBargs *args);
    int s, t;
    DWORD threadID;
    int counter;
    CRITICAL_SECTION counterLock;
    int blockFirst, blockSize, blockRemainder;
    BOOL returnOK;
    DWORD returnCode;

    CHECKINPUTVALUE(numStatements >= 0);
    CHECKINPUTVALUE(statement != NULL);
    for (s = 0; s < numStatements; s++)
        CHECKINPUTVALUE(statement[s] != NULL);
    CHECKINPUTVALUE(ValidMapping(mapping));

```

WO 00/36491

PCT/US99/30274

```

if (mapping != STHREADS_MAPPING_SIMPLE)
    CHECKINPUTVALUE((numThreads > 0) ||
        (numThreads == 0 && numStatements == 0));
CHECKINPUTVALUE(
    ValidPriority(priority) || priority == STHREADS_PRIORITY_PARENT);
CHECKINPUTVALUE(ValidStackSize(stackSize));

if (numStatements == 0) return STHREADS_ERROR_NONE;

if (mapping == STHREADS_MAPPING_SIMPLE) numThreads = numStatements;
if (numThreads > numStatements) numThreads = numStatements;
if (numThreads == 1) mapping = STHREADS_MAPPING_BLOCKED;
if (numThreads == numStatements) mapping = STHREADS_MAPPING_SIMPLE;

CHECKMEMORYALLOC(numThreads <= INT_MAX/sizeof(HANDLE));
thread = (HANDLE *) malloc(numThreads*sizeof(HANDLE));
CHECKMEMORYALLOC(thread != NULL);
CHECKMEMORYALLOC(numThreads <= INT_MAX/sizeof(MTBargs));
threadArgs = (MTBargs *) malloc(numThreads*sizeof(MTBargs));
CHECKMEMORYALLOC(threadArgs != NULL);

parentThread = GetCurrentThread();
believe(parentThread != NULL);
parentPriority = GetThreadPriority(parentThread);
believe(parentPriority != THREAD_PRIORITY_ERROR_RETURN);
believe(ValidPriority(parentPriority));
if (priority != STHREADS_PRIORITY_PARENT) {
    returnOK = SetThreadPriority(parentThread, priority);
    believe(returnOK);
}

switch (mapping) {
case STHREADS_MAPPING_SIMPLE:
    threadStart = SMTBthread;
    break;
case STHREADS_MAPPING_DYNAMIC:
    counter = numThreads - 1;
    InitializeCriticalSection(&counterLock);
    threadStart = DMTBthread;
    break;
case STHREADS_MAPPING_BLOCKED:
    blockFirst = 0;
    blockSize = numStatements/numThreads;
    blockRemainder = numStatements%numThreads;
    threadStart = BMTBthread;
    break;
case STHREADS_MAPPING_INTERLEAVED:
    blockSize = numStatements/numThreads;
    blockRemainder = numStatements%numThreads;
    threadStart = BMTBthread;
    break;
default:
    assert(false);
}

threadCount = numThreads;
threadsFinished = CreateEvent(NULL, TRUE, FALSE, NULL);
CHECKSYNCCREATE(threadsFinished != NULL);
for (t = 0; t < numThreads; t++) {
    threadArgs[t].numStatements = numStatements;
    threadArgs[t].statement = statement;
    threadArgs[t].args = args;
    threadArgs[t].threadCount = (LPLONG) &threadCount;
    threadArgs[t].threadsFinished = threadsFinished;

    switch (mapping) {
    case STHREADS_MAPPING_SIMPLE:
        threadArgs[t].first = t;
        break;
    case STHREADS_MAPPING_DYNAMIC:
        threadArgs[t].first = t;
        threadArgs[t].counter = &counter;
        threadArgs[t].counterLock = &counterLock;
        break;
    }
}

```

WO 00/36491

PCT/US99/30274

```

        case STHREADS_MAPPING_BLOCKED:
            threadArgs[t].first = blockFirst;
            threadArgs[t].last = blockFirst + (blockSize - 1);
            threadArgs[t].step = 1;
            if (blockRemainder > 0) {
                threadArgs[t].last = threadArgs[t].last + 1;
                blockRemainder = blockRemainder - 1;
            }
            blockFirst = threadArgs[t].last + 1;
            break;
        case STHREADS_MAPPING_INTERLEAVED:
            threadArgs[t].first = t;
            threadArgs[t].last = blockSize*numThreads + t;
            threadArgs[t].step = numThreads;
            if (blockRemainder == 0)
                threadArgs[t].last = threadArgs[t].last - numThreads;
            else
                blockRemainder = blockRemainder - 1;
            break;
        default:
            believe(false);
    }

    thread[t] = CreateThread(NULL, stackSize,
        (LPTHREAD_START_ROUTINE) threadStart,
        (LPVOID) &threadArgs[t], CREATE_SUSPENDED, &threadID);
    CHECKTHREADCREATE(thread[t] != NULL);
    if (priority == STHREADS_PRIORITY_PARENT)
        returnOK = SetThreadPriority(thread[t], parentPriority);
    else
        returnOK = SetThreadPriority(thread[t], priority);
    CHECKTHREADCREATE(returnOK);
    returnCode = ResumeThread(thread[t]);
    CHECKTHREADCREATE(returnCode == 1);
}

if (priority != STHREADS_PRIORITY_PARENT) {
    returnOK = SetThreadPriority(parentThread, parentPriority);
    believe(returnOK);
}

returnCode = WaitForSingleObject(threadsFinished, INFINITE);
CHECKOTHER(returnCode != WAIT_FAILED);
returnOK = CloseHandle(threadsFinished);
CHECKOTHER(returnOK == TRUE);
for (t = 0; t < numThreads; t++) {
    returnOK = CloseHandle(thread[t]);
    CHECKOTHER(returnOK == TRUE);
}
if (mapping == STHREADS_MAPPING_DYNAMIC)
    DeleteCriticalSection(&counterLock);
free(thread);
free(threadArgs);

return STHREADS_ERROR_NONE;
}

/*-----*/
/* Is regular for loop range infinite? */
/*-----*/

static bool InfiniteRange(int initial, int condition, int bound, int step)
{
    require(ValidCondition(condition));

    switch (condition) {
        case STHREADS_CONDITION_LT:
            return initial < bound && step <= 0;
        case STHREADS_CONDITION_LE:
            return initial <= bound && step <= 0;
        case STHREADS_CONDITION_GT:
            return initial > bound && step >= 0;
        case STHREADS_CONDITION_GE:
            return initial >= bound && step >= 0;
        default:

```

WO 00/36491

PCT/US99/30274

```

        believe(false);
        return false; /* This return should never be executed. */
    }
}

/*-----*/
/* Is regular for loop range null? */
/*-----*/

static bool NullRange(int initial, int condition, int bound, int step)
{
    require(ValidCondition(condition));

    switch (condition) {
    case STHREADS_CONDITION_LT:
        return initial >= bound;
    case STHREADS_CONDITION_LE:
        return initial > bound;
    case STHREADS_CONDITION_GT:
        return initial <= bound;
    case STHREADS_CONDITION_GE:
        return initial < bound;
    default:
        believe(false);
        return false; /* This return should never be executed. */
    }
}

/*-----*/
/* Arithmetic operations on signed and unsigned integers */
/*-----*/

static unsigned int DIFF(int high, int low)
{
    require(low <= high);

    return (unsigned int) (high - low);
}

/*-----*/

static int ADD(int base, unsigned int offset)
{
    require(offset <= DIFF(INT_MAX, base));

    return base + (int) offset;
}

/*-----*/

static int SUBTRACT(int base, unsigned int offset)
{
    require(offset <= DIFF(base, INT_MIN));

    return base - (int) offset;
}

/*-----*/
/* Split range 0 .. rangeLast into chunks numbered 0 .. chunkLast with */
/* chunks. Return the first and last indices of chunk c. */
/*-----*/

static void SPLIT(
    unsigned int rangeLast, unsigned int chunkLast, unsigned int c,
    unsigned int *first, unsigned int *last)
{
    unsigned int smallerChunkSize;
    unsigned int numLargerChunks;

    require(chunkLast <= rangeLast);
    require(c <= chunkLast);
    require(first != NULL && last != NULL);

    if (chunkLast == 0) {

```

WO 00/36491

PCT/US99/30274

```

        *first = 0;
        *last = rangeLast;
    } else if (chunkLast == rangeLast) {
        *first = c;
        *last = c;
    } else {
        smallerChunkSize = (rangeLast - chunkLast)/(chunkLast + 1) + 1;
        numLargerChunks = (rangeLast - chunkLast)%(chunkLast + 1);
        *first = c*smallerChunkSize + MIN(c, numLargerChunks);
        *last = *first + (smallerChunkSize - 1);
        if (c < numLargerChunks) *last = *last + 1;
    }
}

/*-----*/
/* Last iteration number in regular for loop range */
/* (iterations numbered 0, 1, 2, ...) */
/*-----*/

static unsigned int LAST_ITERATION_NUM(
    int initial, int condition, int bound, int step)
{
    require(ValidCondition(condition));
    require(!InfiniteRange(initial, condition, bound, step));
    require(!NullRange(initial, condition, bound, step));

    switch (condition) {
    case STHREADS_CONDITION_LT:
        believe(initial < bound && step > 0);
        return DIFF(bound - 1, initial)/((unsigned int) step);
    case STHREADS_CONDITION_LE:
        believe(initial <= bound && step > 0);
        return DIFF(bound, initial)/((unsigned int) step);
    case STHREADS_CONDITION_GT:
        believe(initial > bound && step < 0);
        return DIFF(initial, bound + 1)/((unsigned int) -step);
    case STHREADS_CONDITION_GE:
        believe(initial >= bound && step < 0);
        return DIFF(initial, bound)/((unsigned int) -step);
    default:
        assert(false);
        return false; /* This return should never be executed. */
    }
}

/*-----*/
/* Last chunk number in regular for loop range (chunks numbered 0, 1, 2, ...) */
/*-----*/

static unsigned int LAST_CHUNK_NUM(
    int initial, int condition, int bound, int step, int chunkSize)
{
    require(ValidCondition(condition));
    require(!InfiniteRange(initial, condition, bound, step));
    require(!NullRange(initial, condition, bound, step));
    require(chunkSize >= 1);

    return LAST_ITERATION_NUM(initial, condition, bound, step)/
        ((unsigned int) chunkSize);
}

/*-----*/
/* Control value on ith iteration of regular for loop range (i = 0, 1, 2, ...) */
/*-----*/

static int ControlValue(unsigned int i, int initial, int step)
{
    require(step != 0);

    if (step > 0)
        return ADD(initial, i*((unsigned int) step));
    else
        return SUBTRACT(initial, i*((unsigned int) -step));
}

```

WO 00/36491

PCT/US99/30274

```

/*-----*/
/* Does control value lie inside regular for loop range? */
/*-----*/

static bool InRange(
    int controlValue, int initial, int condition, int bound, int step)
{
    require(ValidCondition(condition));
    require(!InfiniteRange(initial, condition, bound, step));
    require(!NullRange(initial, condition, bound, step));

    switch (condition) {
    case STHREADS_CONDITION_LT:
        believe(step > 0);
        return initial <= controlValue && controlValue < bound;
    case STHREADS_CONDITION_LE:
        believe(step > 0);
        return initial <= controlValue && controlValue <= bound;
    case STHREADS_CONDITION_GT:
        believe(step < 0);
        return initial >= controlValue && controlValue > bound;
    case STHREADS_CONDITION_GE:
        believe(step < 0);
        return initial >= controlValue && controlValue >= bound;
    default:
        believe(false);
        return false; /* This return should never be executed. */
    }
}

/*-----*/
/* Execute cth chunk of regular for loop range (c = 0, 1, 2, ...) */
/*-----*/

static void ExecuteChunk(
    int initial, int condition, int bound, int step, int chunkSize,
    unsigned int c, void (*chunk)(int, int, int, void *), void *args)
{
    unsigned int iFirst, iLast;
    int chunkInitial, chunkLast, chunkBound;

    require(ValidCondition(condition));
    require(!InfiniteRange(initial, condition, bound, step));
    require(!NullRange(initial, condition, bound, step));
    require(chunkSize >= 1);
    require(c <= LAST_CHUNK_NUM(initial, condition, bound, step, chunkSize));
    require(chunk != NULL);

    SPLIT(
        LAST_ITERATION_NUM(initial, condition, bound, step),
        LAST_CHUNK_NUM(initial, condition, bound, step, chunkSize), c,
        &iFirst, &iLast);
    believe(0 <= iFirst);
    believe(iFirst <= iLast);
    believe(iLast <= LAST_ITERATION_NUM(initial, condition, bound, step));
    chunkInitial = ControlValue(iFirst, initial, step);
    believe(InRange(chunkInitial, initial, condition, bound, step));
    chunkLast = ControlValue(iLast, initial, step);
    believe(InRange(chunkLast, initial, condition, bound, step));
    switch (condition) {
    case STHREADS_CONDITION_LT:
        chunkBound = chunkLast + 1;
        break;
    case STHREADS_CONDITION_LE:
        chunkBound = chunkLast;
        break;
    case STHREADS_CONDITION_GT:
        chunkBound = chunkLast - 1;
        break;
    case STHREADS_CONDITION_GE:
        chunkBound = chunkLast;
        break;
    default:

```

WO 00/36491

PCT/US99/30274

```

        believe(false);
    }
    (*chunk)(chunkInitial, chunkBound, step, args);
}

/*-----*/
/* Arguments for multithreaded regular for loop thread */
/*-----*/

typedef struct {
    void (*chunk)(int initial, int bound, int step, void *args);
    void *args;
    int initial, condition, bound, step;
    int chunkSize;
    unsigned int chunkFirst, chunkLast, chunkStep;
    unsigned int *counter;
    LPCRITICAL_SECTION counterLock;
    LPLONG threadCount;
    HANDLE threadsFinished;
} MTRFLargs;

/*-----*/
/* Simple multithreaded regular for loop thread */
/*-----*/

static void SMTRFLthread(MTRFLargs *args)
{
    BOOL returnOK;

    require(args != NULL);
    require(args->chunk != NULL);
    require(ValidCondition(args->condition));
    require(!InfiniteRange(
        args->initial, args->condition, args->bound, args->step));
    require(!NullRange(
        args->initial, args->condition, args->bound, args->step));
    require(args->chunkSize >= 1);
    require(args->chunkFirst <= LAST_CHUNK_NUM(
        args->initial, args->condition, args->bound, args->step,
        args->chunkSize));

    ExecuteChunk(
        args->initial, args->condition, args->bound, args->step,
        args->chunkSize, args->chunkFirst, args->chunk, args->args);

    if (InterlockedDecrement(&args->threadCount) == 0) {
        returnOK = SetEvent(args->threadsFinished);
        check(returnOK);
    }
}

/*-----*/
/* Dynamic multithreaded regular for loop thread */
/*-----*/

static void DMTRFLthread(MTRFLargs *args)
{
    unsigned int c, last_c;
    bool finished;
    BOOL returnOK;

    require(args != NULL);
    require(args->chunk != NULL);
    require(ValidCondition(args->condition));
    require(!InfiniteRange(
        args->initial, args->condition, args->bound, args->step));
    require(!NullRange(
        args->initial, args->condition, args->bound, args->step));
    require(args->chunkSize >= 1);
    require(args->chunkFirst <= LAST_CHUNK_NUM(
        args->initial, args->condition, args->bound, args->step,
        args->chunkSize));
    require(args->counter != NULL);
    require(args->counterLock != NULL);

```


WO 00/36491

PCT/US99/30274

```

    c = args->chunkFirst;
    last_c = LAST_CHUNK_NUM(
        args->initial, args->condition, args->bound, args->step,
        args->chunkSize);
    while (true) {
        ExecuteChunk(
            args->initial, args->condition, args->bound, args->step,
            args->chunkSize, c, args->chunk, args->args);
        EnterCriticalSection(args->counterLock);
        finished = (*args->counter == last_c);
        if (!finished) {
            *args->counter = *args->counter + 1;
            c = *args->counter;
        }
        LeaveCriticalSection(args->counterLock);
        if (finished) break;
    }

    if (InterlockedDecrement(args->threadCount) == 0) {
        returnOK = SetEvent(args->threadsFinished);
        check(returnOK);
    }
}

/*-----*/
/* Blocked and interleaved multithreaded regular for loop thread */
/*-----*/

static void BIMTRFLthread(MTRFLargs *args)
{
    unsigned int c;
    BOOL returnOK;

    require(args != NULL);
    require(args->chunk != NULL);
    require(ValidCondition(args->condition));
    require(!InfiniteRange(
        args->initial, args->condition, args->bound, args->step));
    require(!NullRange(
        args->initial, args->condition, args->bound, args->step));
    require(args->chunkSize >= 1);
    require(args->chunkFirst <= args->chunkLast);
    require(args->chunkLast <= LAST_CHUNK_NUM(
        args->initial, args->condition, args->bound, args->step,
        args->chunkSize));
    require((args->chunkLast - args->chunkFirst)%args->chunkStep == 0);

    c = args->chunkFirst;
    while (true) {
        ExecuteChunk(
            args->initial, args->condition, args->bound, args->step,
            args->chunkSize, c, args->chunk, args->args);
        if (c == args->chunkLast) break;
        believe(args->chunkLast - c >= args->chunkStep);
        c = c + args->chunkStep;
    }

    if (InterlockedDecrement(args->threadCount) == 0) {
        returnOK = SetEvent(args->threadsFinished);
        check(returnOK);
    }
}

/*-----*/
/* Multithreaded regular for loop */
/*-----*/

int SthreadsRegularForLoop(
    void (*chunk)(int initial, int bound, int step, void *args), void *args,
    int initial, int condition, int bound, int step,
    int chunkSize, int mapping, int numThreads,
    int priority, unsigned int stackSize)
{

```

WO 00/36491

PCT/US99/30274

```

unsigned int lastChunkNum;
HANDLE *thread;
MTRFLargs *threadArgs;
LONG threadCount;
HANDLE threadsFinished;
HANDLE parentThread;
int parentPriority;
void (*thread_start)(MTRFLargs *args);
int t;
DWORD threadID;
int counter;
CRITICAL_SECTION counterLock;
unsigned int blockFirst, blockSize, blockRemainder;
BOOL returnOK;
DWORD returnCode;

CHECKINPUTVALUE(chunk != NULL);
CHECKINPUTVALUE(ValidCondition(condition));
CHECKINPUTVALUE(!InfiniteRange(initial, condition, bound, step));
CHECKINPUTVALUE((chunkSize > 0) ||
    (chunkSize == 0 &&
        NullRange(initial, condition, bound, step)));
CHECKINPUTVALUE(ValidMapping(mapping));
if (mapping != STHREADS_MAPPING_SIMPLE)
    CHECKINPUTVALUE((numThreads > 0) ||
        (numThreads == 0 &&
            NullRange(initial, condition, bound, step)));
CHECKINPUTVALUE(
    ValidPriority(priority) || priority == STHREADS_PRIORITY_PARENT);
CHECKINPUTVALUE(ValidStackSize(stackSize));

if (NullRange(initial, condition, bound, step))
    return STHREADS_ERROR_NONE;

lastChunkNum = LAST_CHUNK_NUM(
    initial, condition, bound, step, chunkSize);
CHECKMEMORYALLOC(!(mapping == STHREADS_MAPPING_SIMPLE &&
    lastChunkNum >= INT_MAX));

if (mapping == STHREADS_MAPPING_SIMPLE)
    numThreads = (int) (lastChunkNum + 1);
if ((unsigned int) (numThreads - 1) > lastChunkNum)
    numThreads = (int) (lastChunkNum + 1);
if (numThreads == 1)
    mapping = STHREADS_MAPPING_INTERLEAVED;
if ((unsigned int) (numThreads - 1) == lastChunkNum)
    mapping = STHREADS_MAPPING_SIMPLE;

CHECKMEMORYALLOC(numThreads <= INT_MAX/sizeof(HANDLE));
thread = (HANDLE *) malloc(numThreads*sizeof(HANDLE));
CHECKMEMORYALLOC(thread != NULL);
CHECKMEMORYALLOC(numThreads <= INT_MAX/sizeof(MTRFLargs));
threadArgs = (MTRFLargs *) malloc(numThreads*sizeof(MTRFLargs));
CHECKMEMORYALLOC(threadArgs != NULL);

parentThread = GetCurrentThread();
believe(parentThread != NULL);
parentPriority = GetThreadPriority(parentThread);
believe(parentPriority != THREAD_PRIORITY_ERROR_RETURN);
believe(ValidPriority(parentPriority));
if (priority != STHREADS_PRIORITY_PARENT) {
    returnOK= SetThreadPriority(parentThread, priority);
    believe(returnOK);
}

switch (mapping) {
case STHREADS_MAPPING_SIMPLE:
    thread_start = SMTRFLthread;
    break;
case STHREADS_MAPPING_DYNAMIC:
    counter = numThreads - 1;
    InitializeCriticalSection(&counterLock);
    thread_start = DMTRFLthread;
    break;

```

WO 00/36491

PCT/US99/30274

```

case STHREADS_MAPPING_BLOCKEL.
    blockFirst = 0;
    blockSize =
        (lastChunkNum - (((unsigned int) numThreads) - 1))/
        ((unsigned int) numThreads) + 1;
    blockRemainder =
        (lastChunkNum - (((unsigned int) numThreads) - 1))%
        ((unsigned int) numThreads);
    thread_start = BIMTRFLthread;
    break;
case STHREADS_MAPPING_INTERLEAVED:
    blockSize =
        (lastChunkNum - (((unsigned int) numThreads) - 1))/
        ((unsigned int) numThreads) + 1;
    blockRemainder =
        (lastChunkNum - (((unsigned int) numThreads) - 1))%
        ((unsigned int) numThreads);
    thread_start = BIMTRFLthread;
    break;
default:
    assert(false);
)

threadCount = numThreads;
threadsFinished = CreateEvent(NULL, TRUE, FALSE, NULL);
CHECKSYNCCREATE(threadsFinished != NULL);
for (t = 0; t < numThreads; t++) {
    threadArgs[t].chunk = chunk;
    threadArgs[t].args = args;
    threadArgs[t].initial = initial;
    threadArgs[t].condition = condition;
    threadArgs[t].bound = bound;
    threadArgs[t].step = step;
    threadArgs[t].chunkSize = chunkSize;
    threadArgs[t].threadCount = (LPLONG) &threadCount;
    threadArgs[t].threadsFinished = threadsFinished;

    switch (mapping) {
    case STHREADS_MAPPING_SIMPLE:
        threadArgs[t].chunkFirst = t;
        break;
    case STHREADS_MAPPING_DYNAMIC:
        threadArgs[t].chunkFirst = t;
        threadArgs[t].counter = &counter;
        threadArgs[t].counterLock = &counterLock;
        break;
    case STHREADS_MAPPING_BLOCKED:
        threadArgs[t].chunkFirst = blockFirst;
        threadArgs[t].chunkLast = blockFirst + (blockSize - 1);
        threadArgs[t].chunkStep = 1;
        if (blockRemainder > 0) {
            threadArgs[t].chunkLast = threadArgs[t].chunkLast + 1;
            blockRemainder = blockRemainder - 1;
        }
        blockFirst = threadArgs[t].chunkLast + 1;
        break;
    case STHREADS_MAPPING_INTERLEAVED:
        threadArgs[t].chunkFirst = t;
        threadArgs[t].chunkLast =
            blockSize*((unsigned int) numThreads) + t;
        threadArgs[t].chunkStep =
            (unsigned int) numThreads;
        if (blockRemainder == 0)
            threadArgs[t].chunkLast =
                threadArgs[t].chunkLast - ((unsigned int) numThreads);
        else
            blockRemainder = blockRemainder - 1;
        break;
    default:
        believe(false);
    }

    thread[t] = CreateThread(NULL, stackSize,
        (LPTHREAD_START_ROUTINE) thread_start,

```

WO 00/36491

PCT/US99/30274

```

        (LPVOID) &threadArgs, CREATE_SUSPENDED, &threadID);
CHECKTHREADCREATE(thread[t] != NULL);
if (priority == STHREADS_PRIORITY_PARENT)
    SetThreadPriority(thread[t], parentPriority);
else
    SetThreadPriority(thread[t], priority);
ResumeThread(thread[t]);
}

if (priority != STHREADS_PRIORITY_PARENT) {
    SetThreadPriority(parentThread, parentPriority);
    believe(returnOK);
}

returnCode = WaitForSingleObject(threadsFinished, INFINITE);
CHECKOTHER(returnCode != WAIT_FAILED);
returnOK = CloseHandle(threadsFinished);
CHECKOTHER(returnOK == TRUE);
for (t = 0; t < numThreads; t++) {
    returnOK = CloseHandle(thread[t]);
    CHECKOTHER(returnOK == TRUE);
}
if (mapping == STHREADS_MAPPING_DYNAMIC)
    DeleteCriticalSection(&counterLock);
free(thread);
free(threadArgs);

return STHREADS_ERROR_NONE;
}

/*-----*/
/* Multithreaded nested regular for loop (for future release?) */
/*-----*/

int SthreadsNestedRegularForLoop(
    int nesting,
    void (*chunk)(int first[], int last[], int step[], void *args),
    void *args,
    int initial[], int condition[], int bound[], int step[],
    int chunkSize[], int mapping[], int numThreads[],
    int priority, unsigned int stackSize)
/* Arguments:
/* - nesting      : degree of nesting.
/* - chunk        : function to execute chunk of iterations of loop body.
/* - args         : pointer to arguments of loop body.
/* - initial      : initial value of control variable at each nesting level.
/* - condition    : condition between control variable and bound value
/*                  at each nesting level.
/* - bound        : bound value of control variable at each nesting level.
/* - step         : step value of control variable at each nesting level.
/* - chunkSize    : number of iterations per chunk at each nesting level.
/* - mapping      : mapping of chunks onto threads at each nesting level.
/* - numThreads   : number of threads at each nesting level.
/* - priority     : priority of threads.
/* - stackSize    : stack size of threads.
/* Returns:
/* - error code.
/* Requirements:
/* - nesting >= 1
/* - chunk != NULL &&
/*   chunk is a valid void (*)(int *, int *, int *, void *) function.
/* - initial != NULL &&
/*   initial is an array of at least nesting ints.
/* - condition != NULL &&
/*   condition is an array of at least nesting ints.
/* - forall (i = 0; i < nesting; i++) ValidCondition(condition[i]).
/* - bound != NULL &&
/*   bound is an array of at least nesting ints.
/* - step != NULL &&
/*   step is an array of at least nesting ints.
/* - forall (i = 0; i < nesting; i++)
/*   !InfiniteRange(initial[i], condition[i], bound[i], step[i]) ||
/*     exists (j = 0; j < i; j++)
/*       NullRange(initial[j], condition[j], bound[j], step[j]).
/* - forall (i = 0; i < nesting; i++)

```

WO 00/36491

PCT/US99/30274

```

/*
    (chunkSize[i] > 0) ||
    (chunkSize[i] == 0 &&
    NullRange(initial[i], condition[i], bound[i], step[i])).
    - forall (i = 0; i < nesting; i++) ValidMapping(mapping[i]).
    - forall (i = 0; i < nesting; i++)
    mapping[i] != STHEADS_MAPPING_SIMPLE =>
    (numThreads[i] > 0) ||
    (numThreads[i] == 0 &&
    NullRange(initial[i], condition[i], bound[i], step[i])).
    - ValidPriority(priority) || priority == STHEADS_PRIORITY_PARENT.
    - ValidStackSize(stackSize).
*/
{
    int i;

    CHECKINPUTVALUE(nesting >= 1);
    CHECKINPUTVALUE(chunk != NULL);
    CHECKINPUTVALUE(initial != NULL);
    CHECKINPUTVALUE(condition != NULL);
    for (i = 0; i < nesting; i++)
        CHECKINPUTVALUE(ValidCondition(condition[i]));
    CHECKINPUTVALUE(bound != NULL);
    CHECKINPUTVALUE(step != NULL);
    for (i = 0; i < nesting; i++) {
        if (NullRange(initial[i], condition[i], bound[i], step[i])) break;
        CHECKINPUTVALUE(
            !InfiniteRange(initial[i], condition[i], bound[i], step[i]));
    }
    for (i = 0; i < nesting; i++)
        CHECKINPUTVALUE((chunkSize[i] > 0) ||
            (chunkSize[i] == 0 &&
            NullRange(initial[i], condition[i], bound[i], step[i])));
    for (i = 0; i < nesting; i++)
        CHECKINPUTVALUE(ValidMapping(mapping[i]));
    for (i = 0; i < nesting; i++)
        if (mapping[i] != STHEADS_MAPPING_SIMPLE)
            CHECKINPUTVALUE(
                (numThreads[i] > 0) ||
                (numThreads[i] == 0 &&
                NullRange(initial[i], condition[i], bound[i], step[i])));
    CHECKINPUTVALUE(
        ValidPriority(priority) || priority == STHEADS_PRIORITY_PARENT);
    CHECKINPUTVALUE(ValidStackSize(stackSize));

    return STHEADS_ERROR_NONE;
}

/*-----*/
/* Multithreaded general for loop (for future release?) */
/*-----*/

int StthreadsGeneralForLoop(
    void (*body)(void *control, void *args),
    size_t controlSize, void *args,
    int (*test)(void *args), void (*increment)(void *args),
    void (*copy)(void *control, void *args),
    int mapping, int numThreads,
    int priority, unsigned int stackSize)
/* Arguments:
/* - body      : function to execute one iteration of loop body.
/* - controlSize : size (as returned by sizeof) of control variables.
/* - args      : pointer to arguments of loop.
/* - test      : function to test loop termination condition.
/* - increment  : function to increment control variables within arguments.
/* - copy      : function to copy control variables from arguments.
/* - mapping    : mapping of iterations onto threads.
/* - numThreads : number of threads.
/* - priority   : priority of threads.
/* - stackSize  : stack size of threads.
/* Returns:
/* - error code.
/* Requirements:
/* - body != NULL &&
/*   body is a valid void (*) (void *, void *) function.
/* - test != NULL &&

```

WO 00/36491

PCT/US99/30274

```

/* test is a valid int (*)(void *) function.
/* - increment != NULL &&
/* increment is a valid void (*)(void *) function.
/* - copy != NULL &&
/* copy is a valid void (*)(void *, void *) function.
/* - mapping == STHREADS_MAPPING_SIMPLE ||
/* mapping == STHREADS_MAPPING_DYNAMIC.
/* - mapping != STHREADS_MAPPING_SIMPLE =>
/* (numThreads > 0) || (numThreads == 0 && !test(args)).
/* - ValidPriority(priority) || priority == STHREADS_PRIORITY_PARENT.
/* - ValidStackSize(stackSize).
{
    CHECKINPUTVALUE(body != NULL);
    CHECKINPUTVALUE(test != NULL);
    CHECKINPUTVALUE(increment != NULL);
    CHECKINPUTVALUE(copy != NULL);
    CHECKINPUTVALUE(mapping == STHREADS_MAPPING_SIMPLE ||
        mapping == STHREADS_MAPPING_DYNAMIC);
    if (mapping != STHREADS_MAPPING_SIMPLE)
        CHECKINPUTVALUE((numThreads > 0) || (numThreads == 0 && !test(args)));
    CHECKINPUTVALUE(
        ValidPriority(priority) || priority == STHREADS_PRIORITY_PARENT);
    CHECKINPUTVALUE(ValidStackSize(stackSize));

    return STHREADS_ERROR_NONE;
}

/*-----*/
/* Synchronization object status constants
/*-----*/

#define INITIALIZED 123456
#define FINALIZED 654321

/*-----*/
/* Flags
/*-----*/

typedef struct {
    int initialized, finalized;
    LONG numWaiting;
    HANDLE signal;
} PrivateFlag;

#define PRIVATE(flagPtr) ((PrivateFlag *) (flagPtr))

/*-----*/

int SthreadsFlagInitialize(SthreadsFlag *flag)
{
    CHECKINPUTVALUE(flag != NULL);

    PRIVATE(flag)->initialized = INITIALIZED;
    PRIVATE(flag)->finalized = -FINALIZED;
    PRIVATE(flag)->numWaiting = 0;
    PRIVATE(flag)->signal = CreateEvent(NULL, TRUE, FALSE, NULL);
    CHECKSYNCCREATE(PRIVATE(flag)->signal != NULL);

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsFlagFinalize(SthreadsFlag *flag)
{
    BOOL returnOK;

    CHECKINPUTVALUE(flag != NULL);
    CHECKUNINITIALIZED(PRIVATE(flag)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(flag)->finalized == -FINALIZED);
    CHECKINUSE(PRIVATE(flag)->numWaiting == 0);

    PRIVATE(flag)->finalized = FINALIZED;
    returnOK = CloseHandle(PRIVATE(flag)->signal);

```

WO 00/36491

PCT/US99/30274

```

        CHECKOTHER(returnOK == TRUE),
        return STHREADS_ERROR_NONE;
    }

/*-----*/
int SthreadsFlagSet(SthreadsFlag *flag)
{
    BOOL returnOK;

    CHECKINPUTVALUE(flag != NULL);
    CHECKUNINITIALIZED(PRIVATE(flag)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(flag)->finalized == -FINALIZED);

    returnOK = SetEvent(PRIVATE(flag)->signal);
    CHECKOTHER(returnOK);

    return STHREADS_ERROR_NONE;
}

/*-----*/
int SthreadsFlagCheck(SthreadsFlag *flag)
{
    DWORD returnCode;

    CHECKINPUTVALUE(flag != NULL);
    CHECKUNINITIALIZED(PRIVATE(flag)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(flag)->finalized == -FINALIZED);

    InterlockedIncrement(&PRIVATE(flag)->numWaiting);
    returnCode = WaitForSingleObject(PRIVATE(flag)->signal, INFINITE);
    CHECKOTHER(returnCode != WAIT_FAILED);
    InterlockedDecrement(&PRIVATE(flag)->numWaiting);

    return STHREADS_ERROR_NONE;
}

/*-----*/
int SthreadsFlagReset(SthreadsFlag *flag)
{
    BOOL returnOK;

    CHECKINPUTVALUE(flag != NULL);
    CHECKUNINITIALIZED(PRIVATE(flag)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(flag)->finalized == -FINALIZED);
    CHECKINUSE(PRIVATE(flag)->numWaiting == 0);

    PRIVATE(flag)->numWaiting = 0;
    returnOK = ResetEvent(PRIVATE(flag)->signal);
    CHECKOTHER(returnOK);

    return STHREADS_ERROR_NONE;
}

/*-----*/
#undef PRIVATE

/*-----*/
/* Counters */
/*-----*/

typedef struct node *link;
typedef struct node {
    unsigned int value;
    int numWaiting;
    HANDLE signal;
    link next;
} node;

typedef struct {

```

WO 00/36491

PCT/US99/30274

```

    int initialized, finalized;
    unsigned int count;
    link waitingList;
    CRITICAL_SECTION lock;
} PrivateCounter;

#define PRIVATE(counterPtr) ((PrivateCounter *) (counterPtr))

/*-----*/

int SthreadsCounterInitialize(SthreadsCounter *counter)
{
    link startSentinel, endSentinel;

    CHECKINPUTVALUE(counter != NULL);

    PRIVATE(counter)->initialized = INITIALIZED;
    PRIVATE(counter)->finalized = -FINALIZED;
    PRIVATE(counter)->count = 0;
    startSentinel = (link) malloc(sizeof(node));
    CHECKMEMORYALLOC(startSentinel != NULL);
    endSentinel = (link) malloc(sizeof(node));
    CHECKMEMORYALLOC(endSentinel != NULL);
    startSentinel->signal = NULL;
    startSentinel->next = endSentinel;
    startSentinel->numWaiting = 0;
    endSentinel->signal = NULL;
    endSentinel->next = NULL;
    endSentinel->numWaiting = 0;
    PRIVATE(counter)->waitingList = startSentinel;
    InitializeCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);

    return STTHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsCounterFinalize(SthreadsCounter *counter)
{
    link p, next;
    BOOL returnOK;

    CHECKINPUTVALUE(counter != NULL);
    CHECKUNINITIALIZED(PRIVATE(counter)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(counter)->finalized == -FINALIZED);
    CHECKINUSE(PRIVATE(counter)->waitingList->next->next == NULL);

    PRIVATE(counter)->finalized = FINALIZED;
    p = PRIVATE(counter)->waitingList;
    next = p->next;
    free(p);
    p = next;
    while (p->next != NULL) {
        returnOK = CloseHandle(p->signal);
        CHECKOTHER(returnOK == TRUE);
        next = p->next;
        free(p);
        p = next;
    }
    free(p);
    DeleteCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);

    return STTHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsCounterIncrement(SthreadsCounter *counter, unsigned int amount)
{
    link start, p;
    BOOL returnOK;

    CHECKINPUTVALUE(counter != NULL);
    CHECKUNINITIALIZED(PRIVATE(counter)->initialized == INITIALIZED);

```


WO 00/36491

PCT/US99/30274

```

CHECKFINALIZED(PRIVATE(counter)->initialized == FINALIZED);
CHECKCOUNTEROVERFLOW(PRIVATE(counter)->count <= UINT_MAX - amount);

EnterCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);
PRIVATE(counter)->count = PRIVATE(counter)->count + amount;
start = PRIVATE(counter)->waitingList;
p = start->next;
while (p->next != NULL && p->value <= PRIVATE(counter)->count) {
    returnOK = SetEvent(p->signal);
    CHECKOTHER(returnOK);
    start->next = p->next;
    p = start->next;
}
LeaveCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);

return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsCounterCheck(SthreadsCounter *counter, unsigned int value)
{
    link prev, p;
    link waitingNode;
    BOOL returnOK;
    DWORD returnCode;

    CHECKINPUTVALUE(counter != NULL);
    CHECKUNINITIALIZED(PRIVATE(counter)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(counter)->finalized == -FINALIZED);

    EnterCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);
    if (PRIVATE(counter)->count >= value)
        LeaveCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);
    else {
        prev = PRIVATE(counter)->waitingList;
        p = prev->next;
        while (p->next != NULL && p->value < value) {
            prev = p;
            p = p->next;
        }
        if (p->value == value) {
            waitingNode = p;
            waitingNode->numWaiting = waitingNode->numWaiting + 1;
        } else {
            waitingNode = (link) malloc(sizeof(node));
            waitingNode->value = value;
            waitingNode->signal = CreateEvent(NULL, TRUE, FALSE, NULL);
            waitingNode->next = p;
            waitingNode->numWaiting = 1;
            prev->next = waitingNode;
        }
        LeaveCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);
        returnCode = WaitForSingleObject(waitingNode->signal, INFINITE);
        CHECKOTHER(returnCode != WAIT_FAILED);
        EnterCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);
        waitingNode->numWaiting = waitingNode->numWaiting - 1;
        if (waitingNode->numWaiting == 0) {
            returnOK = CloseHandle(waitingNode->signal);
            CHECKOTHER(returnOK == TRUE);
            free(waitingNode);
        }
        LeaveCriticalSection((LPCRITICAL_SECTION) &PRIVATE(counter)->lock);
    }

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsCounterReset(SthreadsCounter *counter)
{
    link p, q;
    BOOL returnOK;

```

WO 00/36491

PCT/US99/30274

```

CHECKINPUTVALUE(counter != NULL);
CHECKUNINITIALIZED(PRIVATE(counter)->initialized == INITIALIZED);
CHECKFINALIZED(PRIVATE(counter)->finalized == ~FINALIZED);
CHECKINUSE(PRIVATE(counter)->waitingList->next->next == NULL);

PRIVATE(counter)->count = 0;
p = PRIVATE(counter)->waitingList;
q = p->next;
while (q->next != NULL) {
    p->next = q->next;
    returnOK = CloseHandle(q->signal);
    CHECKOTHER(returnOK == TRUE);
    free(q);
    q = p->next;
}

return STTHREADS_ERROR_NONE;
}

/*-----*/
#undef PRIVATE

/*-----*/
/* Locks */
/*-----*/

typedef struct {
    int initialized, finalized;
    HANDLE holder;
    CRITICAL_SECTION lock;
} PrivateLock;

#define PRIVATE(lockPtr) ((PrivateLock *) (lockPtr))

/*-----*/

int SthreadsLockInitialize(SthreadsLock *lock)
{
    CHECKINPUTVALUE(lock != NULL);

    PRIVATE(lock)->initialized = INITIALIZED;
    PRIVATE(lock)->finalized = ~FINALIZED;
    PRIVATE(lock)->holder = NULL;
    InitializeCriticalSection((LPCRITICAL_SECTION) &PRIVATE(lock)->lock);

    return STTHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsLockFinalize(SthreadsLock *lock)
{
    CHECKINPUTVALUE(lock != NULL);
    CHECKUNINITIALIZED(PRIVATE(lock)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(lock)->finalized == ~FINALIZED);
    CHECKINUSE(PRIVATE(lock)->holder == NULL);

    PRIVATE(lock)->finalized = FINALIZED;
    DeleteCriticalSection((LPCRITICAL_SECTION) &PRIVATE(lock)->lock);

    return STTHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsLockAcquire(SthreadsLock *lock)
{
    HANDLE thisThread;

    thisThread = GetCurrentThread();
    believe(thisThread != NULL);

```

WO 00/36491

PCT/US99/30274

```

int SthreadsBarrierFinalize(SthreadsBarrier *barrier)
{
    BOOL returnOK;

    CHECKINPUTVALUE(barrier != NULL);
    CHECKUNINITIALIZED(PRIVATE(barrier)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(barrier)->finalized == -FINALIZED);
    CHECKINUSE(PRIVATE(barrier)->numWaiting == 0);

    PRIVATE(barrier)->finalized = FINALIZED;
    returnOK = CloseHandle(PRIVATE(barrier)->gate[0]);
    CHECKOTHER(returnOK == TRUE);
    returnOK = CloseHandle(PRIVATE(barrier)->gate[1]);
    CHECKOTHER(returnOK == TRUE);
    DeleteCriticalSection((LPCRITICAL_SECTION) &PRIVATE(barrier)->lock);

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsBarrierPass(SthreadsBarrier *barrier)
{
    int currentGate, nextGate;
    BOOL returnOK;
    DWORD returnCode;

    CHECKINPUTVALUE(barrier != NULL);
    CHECKUNINITIALIZED(PRIVATE(barrier)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(barrier)->finalized == -FINALIZED);

    EnterCriticalSection((LPCRITICAL_SECTION) &PRIVATE(barrier)->lock);
    currentGate = PRIVATE(barrier)->currentGate;
    PRIVATE(barrier)->numWaiting = PRIVATE(barrier)->numWaiting + 1;
    if (PRIVATE(barrier)->numWaiting == PRIVATE(barrier)->numThreads) {
        nextGate = (currentGate + 1) % 2;
        returnOK = ResetEvent(PRIVATE(barrier)->gate[nextGate]);
        CHECKOTHER(returnOK);
        PRIVATE(barrier)->numWaiting = 0;
        returnOK = SetEvent(PRIVATE(barrier)->gate[currentGate]);
        CHECKOTHER(returnOK);
        PRIVATE(barrier)->currentGate = nextGate;
        LeaveCriticalSection((LPCRITICAL_SECTION) &PRIVATE(barrier)->lock);
    } else {
        LeaveCriticalSection((LPCRITICAL_SECTION) &PRIVATE(barrier)->lock);
        returnCode = WaitForSingleObject(
            PRIVATE(barrier)->gate[currentGate], INFINITE);
        CHECKOTHER(returnCode != WAIT_FAILED);
    }

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsBarrierReset(SthreadsBarrier *barrier, int numThreads)
{
    BOOL returnOK;

    CHECKINPUTVALUE(barrier != NULL);
    CHECKUNINITIALIZED(PRIVATE(barrier)->initialized == INITIALIZED);
    CHECKFINALIZED(PRIVATE(barrier)->finalized == -FINALIZED);
    CHECKINUSE(PRIVATE(barrier)->numWaiting == 0);
    CHECKINPUTVALUE(numThreads >= 1);

    PRIVATE(barrier)->numThreads = numThreads;
    PRIVATE(barrier)->numWaiting = 0;
    returnOK = ResetEvent(PRIVATE(barrier)->gate[0]);
    CHECKOTHER(returnOK);
    returnOK = SetEvent(PRIVATE(barrier)->gate[1]);
    CHECKOTHER(returnOK);
    PRIVATE(barrier)->currentGate = 0;
}

```

WO 00/36491

PCT/US99/30274

```
-
    return STHREADS_ERROR_NONE;
}

/*-----*/

#undef PRIVATE

/*-----*/
/* Priorities */
/*-----*/

int SthreadsGetCurrentPriority(int *priority)
{
    HANDLE currentThread;
    int currentPriority;

    CHECKINPUTVALUE(priority != NULL);

    currentThread = GetCurrentThread();
    believe(currentThread != NULL);
    currentPriority = GetThreadPriority(currentThread);
    believe(currentPriority != THREAD_PRIORITY_ERROR_RETURN);

    *priority = currentPriority;

    return STHREADS_ERROR_NONE;
}

/*-----*/

int SthreadsSetCurrentPriority(int priority)
{
    HANDLE currentThread;
    BOOL returnOK;

    CHECKINPUTVALUE(ValidPriority(priority));

    currentThread = GetCurrentThread();
    believe(currentThread != NULL);
    returnOK = SetThreadPriority(currentThread, priority);
    believe(returnOK);

    return STHREADS_ERROR_NONE;
}

/*-----*/
```

WO 00/36491

PCT/US99/30274

What is claimed is:

1. A method of synchronizing threads in a multiple
5 thread system, comprising:
 defining an entity which maintains a count of values
 which increases the value maintained by the object; and
 defining a check operation for said element in
 which, during the checking operation, a calling thread is
10 suspended, and the check is suspended until the value
 maintained by the entity has reached or exceeded a given
 value.
2. A method as in claim 1 which said entity is
 allowed only to increment between allowable values, and
15 not to decrement its value.
3. A method as in claim 1 wherein said entity is a
 counter that is only allowed to include integers.
4. A method as in claim 3 wherein an initial value of
 the counter is zero.
- 20 5. A method as in claim 1, wherein said
 entity is a/are flags.
6. An apparatus comprising a machine-readable
 storage medium having executable instructions for
 managing threads in a multithreaded system, the
25 instructions enabling the machine to:

WO 00/36491

PCT/US99/30274

define an entity which maintains a count of values
and which is allowed to increment between allowable
values;

determine a request for value of the element from a
5 calling thread; and

establish a check operation for said element in
which said calling thread is suspended until the entity
reaches a predetermined value.

7. An apparatus as in claim 6, wherein said entity is a
10 monotonically increasing counter.

8. An apparatus as in claim 6, wherein said entity
is a flag.

9. A apparatus as in claim 6 wherein said system
has a plurality of processors therein, wherein each of
15 said processors is running at least one different ones of
said threads.

10. A method as in claim 1, further comprising
defining an error for an operation that decreases the
value maintained by the object to occur concurrently with
20 any check operation on the object.

11. A method as in claim 1, wherein the value
maintained by the object is a numeric value and the
increment operation increases the value by a numeric
amount.

WO 00/36491

PCT/US99/30274

12. A method as in claim 1, wherein the value maintained by the object is a Boolean value or a binary value and the increment operation is a "set" operation that changes the value from one state to the other state.

5 13. A method as in claim 2, wherein the value maintained by the object is a Boolean value or a binary value and the increment operation is a "set" operation that changes the value from one state to the other state.

10 14. A method as in claim 12, further comprising establishing an error for an increment operation on the object to occur more than once.

 15. A method of defining program code, comprising:
 determining different parts of a program which can
15 be executed either sequentially, or in multithreaded parallel by different threads, and which has equivalent results when executed in said sequential or multithreaded parallel; and

 defining said different parts as being multi-
20 threadable.

 16. A method as in claim 15 wherein said determining is based on a set of conditions that are sufficient to ensure the equivalence of sequential and multithreaded execution of a program construct.

WO 00/36491

PCT/US99/30274

17. A method as in claim 15 wherein said different parts are defined as being multithreadable using an equivalence annotation within the program code.

18. A method as in claim 17 wherein said annotation
5 is a pragma.

19. A method as in claim 17 wherein said annotation is a code comment.

20. A method as in claim 15 further comprising,
within said code, multithreaded constructs, in addition
10 to said multithreadable parts.

21. A method as in claim 15 wherein said multithreadable parts includes information which, if executed as threads, will include the same result as if executed sequentially.

15 22. A method as in claim 15 wherein said part is a multithreadable block of information.

23. A method as in claim 22 wherein said part is a multithreadable for loop.

24. A method as in claim 15 further comprising
20 synchronizing threads using a monotonically-increasing counter.

25. A method as in claim 15 further comprising synchronizing threads using a flag.

WO 00/36491

PCT/US99/30274

26. A method as in claim 16, wherein the equivalence annotation includes a new or existing keyword or reserved word in the program.

27. A method as in claim 16, wherein the
5 equivalence annotation takes the form of a character formatting in the program, which can be such as boldface, italics, underlining, or other formatting.

28. A method as in claim 16, wherein the
10 equivalence annotation takes the form of a special character sequence in the program.

29. A method as in claim 16, wherein the equivalence annotation is contained in a file or other entity separate from the program.

30. A method as in claim 16, wherein the sequential
15 interpretation of the execution of the block construct is that statements are executed one at a time in their textual order, and the multithreaded interpretation of the execution of the block construct is that statements of are partitioned among a set of threads and executed
20 concurrently by those threads.

31. A method as in claim 16 further comprising using monotonic thread synchronization to synchronize actions among threads.

WO 00/36491

PCT/US99/30274

32. A method as in claim 15 wherein:

explicitly multithreaded program constructs are
always executed according to a multithreaded
interpretation

5 multithreadable program constructs are either
executed according a multithreaded interpretation or
executed according to a sequential interpretation; and
sequential or multithreaded execution of
multithreadable program constructs is at user selection.

10 33. A method as in claim 32, wherein the sequential
or multithreaded execution of multithreadable program
constructs is signalled by a pragma in the program.

34. A method as in claim 32, wherein the method for
selecting sequential or multithreaded execution of
15 multithreadable code constructs is a variable that is
dependent of the value of a variable defined in the
program or in the environment of the program.

35. A method of claim 32 wherein said multiple
threaded construct is a block or for loop.

20 36. A method of coding a program, comprising:
defining a first portion of code which must always
be executed according to multithreaded semantics, as a
multithreaded portion of code;

WO 00/36491

PCT/US99/30274

defining a second portion of code, within the same
program as said first portion of code, which may be
selectively executed according to either sequential or
multithreaded techniques, as a multithreadable code
5 construct; and

allowing a program development system to develop
said multithreadable code construct as either a
sequential or multithreaded construct.

37. A method as in claim 36, wherein said program
10 development system includes a compiler.

38. A method as in claim 36 wherein said
multithreaded construct defines an operation which has no
sequential equivalent.

39. A method as in claim 38 wherein said
15 multithreaded construct is control of multiple windows in
a graphical system.

40. A method as in claim 38 wherein said
multithreaded construct is control of different
operations of a computer.

20 41. A method as in claim 37 wherein said operation
is executed on a multiple processor system, and different
parts of said operation are executed on different ones of
the processors.

WO 00/36491

PCT/US99/30274

42. A method as in claim 37 wherein said multithreadable constructs include a synchronization mechanism.

43. A method as in claim 42 wherein said
5 synchronization mechanism is a monotonically increasing counter.

44. A method as in claim 43 wherein said synchronization mechanism is a special flag.

45. A method of integrating a structured
10 multithreading program development system with a standard program development system, comprising:

detecting program elements which include a specified annotation;

calling a special program development system element
15 which includes a processor that modifies based on the annotation to form a preprocessed file; and

calling the standard program development system to compile the preprocessed file.

46. A method of operating a program language,
20 comprising:

defining equivalence annotations within the programming language which indicate to a program development system of the programming language

WO 00/36491

PCT/US99/30274

information about sequential execution of said statement;
and

developing the programs as a sequential execution or
as a substantially simultaneous execution based on
5 contents of the equivalence annotations.

47. A method as in claim 46 wherein the equivalence
annotation indicates that the statements are
multithreadable.

48. A method as in claim 46 wherein the equivalence
10 annotation indicates that the statements are either
multithreaded or multithreadable.

49. A method as in claim 48 wherein said
multithreaded statements must be executed in a
multithreaded manner.

15 50. A method as in claim 48 wherein said
multithreadable annotations indicate that the statements
can be executed in either multithreaded or sequential
manner.

51. A method as in claim 46 wherein said
20 equivalence annotation is a pragma.

52. A method as in claim 46 wherein said
equivalence annotation is a specially-defined comment
line.

WO 00/36491

PCT/US99/30274

53. A method as in claim 47 further comprising
synchronizing access of threads to shared memory using a
specially defined synchronization element.

54. A method as in claim 53 wherein said
5 synchronization element is a synchronization counter.

55. A method as in claim 54 wherein said
synchronization counter is monotonically increasing,
cannot be decreased, and prevents thread operation during
its check operation.

10 56. A method as in claim 53 wherein said
synchronization element is a synchronization flag.

57. A method as in claim 56 wherein said
synchronization counter is monotonically increasing,
cannot be decreased, and prevents thread operation during
15 its check operation.

58. A method as in claim 54 wherein said s counter
includes a check operation, wherein said check operation
suspends a calling thread.

59. A method as in claim 58 further comprising
20 maintaining a list of suspended threads.

WO 00/36491

PCT/US99/30274

60. A method of modifying an existing program development system and environment, comprising:

- detecting which components of a program contain multithreadable program constructs or explicitly

5 multithreaded program constructs;

- transforming the components of the program that contain multithreadable program constructs or explicitly multithreaded program constructs into equivalent multithreaded components in a form that can be directly

10 translated or executed by the existing program development system; and

invoking the existing program development system to translate or execute the transformed components of the program.

15 61. A method as in claim 60, wherein said indicating comprises giving distinctive names to said component.

62. A method as in claim 59, wherein the transforming of the components of the program that
20 contain multithreadable program constructs or explicitly multithreaded program constructs is by source-to-source program preprocessing.

63. A method as in claim 61, wherein the result of the source-to-source program preprocessing is a program

WO 00/36491

PCT/US99/30274

component that incorporates thread library calls
representing to the transformed multithreadable program
constructs or explicitly multithreaded program
constructs.

5 64. A method as in claim 63, wherein the thread
library is a thread library designed in part or whole for
the purpose of representing the transformed
multithreadable program constructs or explicitly
multithreaded program constructs.

10 65. A method as in claim 63, wherein the thread
library is an existing thread library or a thread library
designed for another purpose.

 66. A method as in claim 61, wherein the result of
the source-to-source program preprocessing is a program
15 component that incorporates standard multithreaded
program constructs supported by the existing programming
system.

 67. A method as in claim 59, further comprising
renaming the standard compiler-linker and the standard
20 compiler-linker name is used for a program component
transformation tool that subsequently invokes the renamed
standard compiler-linker.

 68. A method as in claim 59, wherein the operating
system is Linux or another variant of the Unix operating

WO 00/36491

PCT/US99/30274

system and the existing program development environment is the GNU C or C++ compiler or any other C or C++ compiler that operates under the given variant of the Unix operating system.

5 69. A method as in claim 59, wherein the existing programming language is a variant of the Java programming language and the thread library is the standard Java thread library.

10 70. A method of operating a program operation, comprising:

 defining a block of code which can be executed either sequentially or substantially simultaneously via separate loci of execution;

 running the program during a first mode in said sequential mode, and running the program during a second mode in said substantially simultaneous mode.

 71. A method as in claim 70 wherein said definition is an equivalence annotation.

20 72. A method as in claim 71 wherein said equivalence annotation is a pragma.

 73. A method as in claim 70 wherein, during said sequential execution, variables are shared.

WO 00/36491

PCT/US99/30274

74. A method as in claim 73 wherein said shared variables can be checked, and operation of check does not suspend operations of the program.

75. A method as in claim 70 wherein during said
5 substantially simultaneous operations, variables are shared.

76. A method as in claim 70 further comprising debugging a program in said sequential mode and running a debugged program in said substantially simultaneous mode.

10 77. An object for synchronizing among multiple threads, comprising:

a special object constrained to have (1) an integer attribute value, (2) an increment function, but no decrement function, and (3) check function that suspends
15 a calling thread.

78. A method as in claim 77 wherein said check function suspends a calling thread for a specified time.

79. An object as in claim 78 wherein said object includes a list of thread suspension queues.

20 80. An object as in claim 77 further comprising a reset function.

81. An object as in claim 77 wherein said object is a counter.

WO 00/36491

PCT/US99/30274

82. An object as in claim 77 wherein said object is a flag having only first and second values.

83. A method of integrating a thread management system with an existing program development system,
5 comprising:

first, running a pre-program development system that looks for special annotations which indicate multithreaded and multithreadable block of code;

using said special layer as an initial linker; and
10 then, passing the already linked program to the standard program development system.

84. A method as in claim 83 wherein said program is a C programming language.

WO 00/36491

PCT/US99/30274

1/4





FIG. 1
(Prior Art)

WO 00/36491

PCT/US99/30274

2/4

F162



Thread Execution Model: Key Points

- Pool of processors, pool of threads.
- Threads are peers.
- Dynamic thread creation.
- Can support many more threads than processors.
- Threads dynamically switch between processors.
- Threads share access to memory.
- Synchronization needed between threads.

November 9th, 1998

SC98 Tutorial, Copyright (c) 1998 by John Thornley.

18

WO 00/36491

PCT/US99/30274

3/4

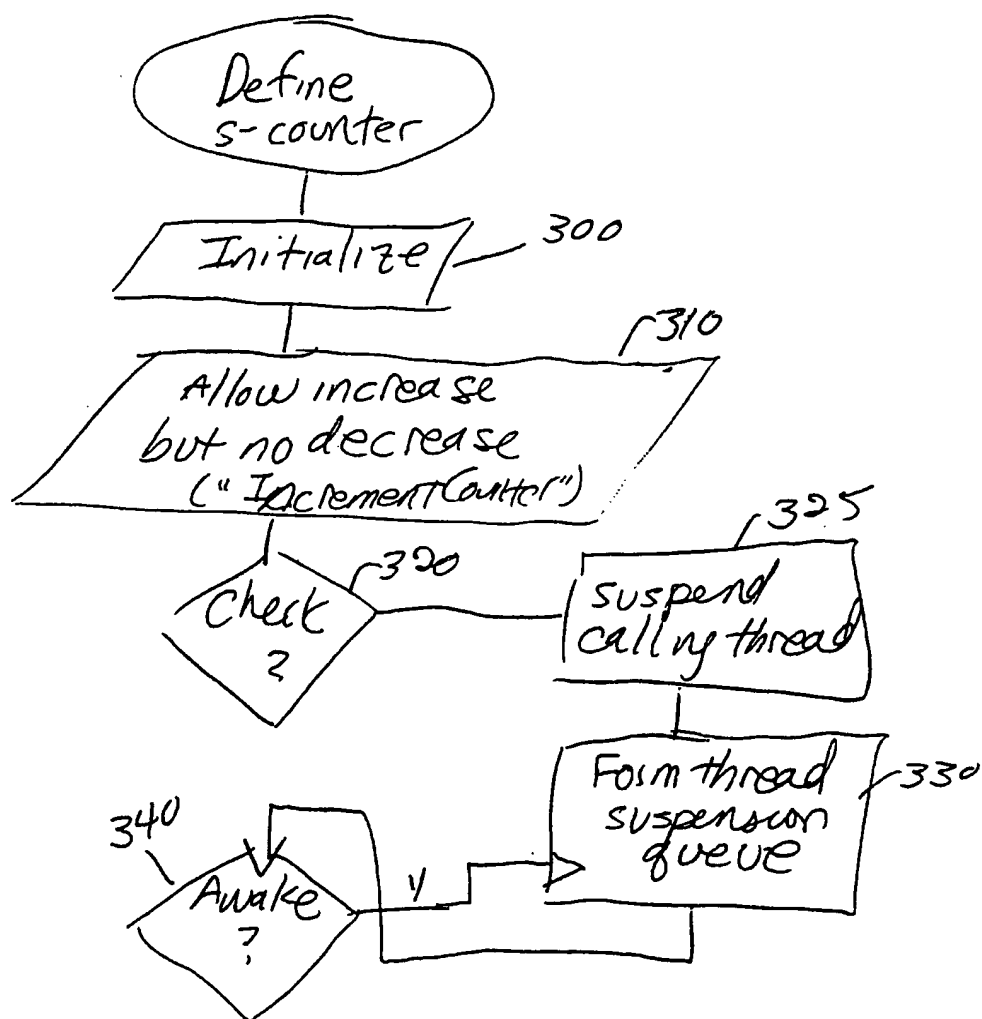


FIG 3

WO 00/36491

PCT/US99/30274

4/4

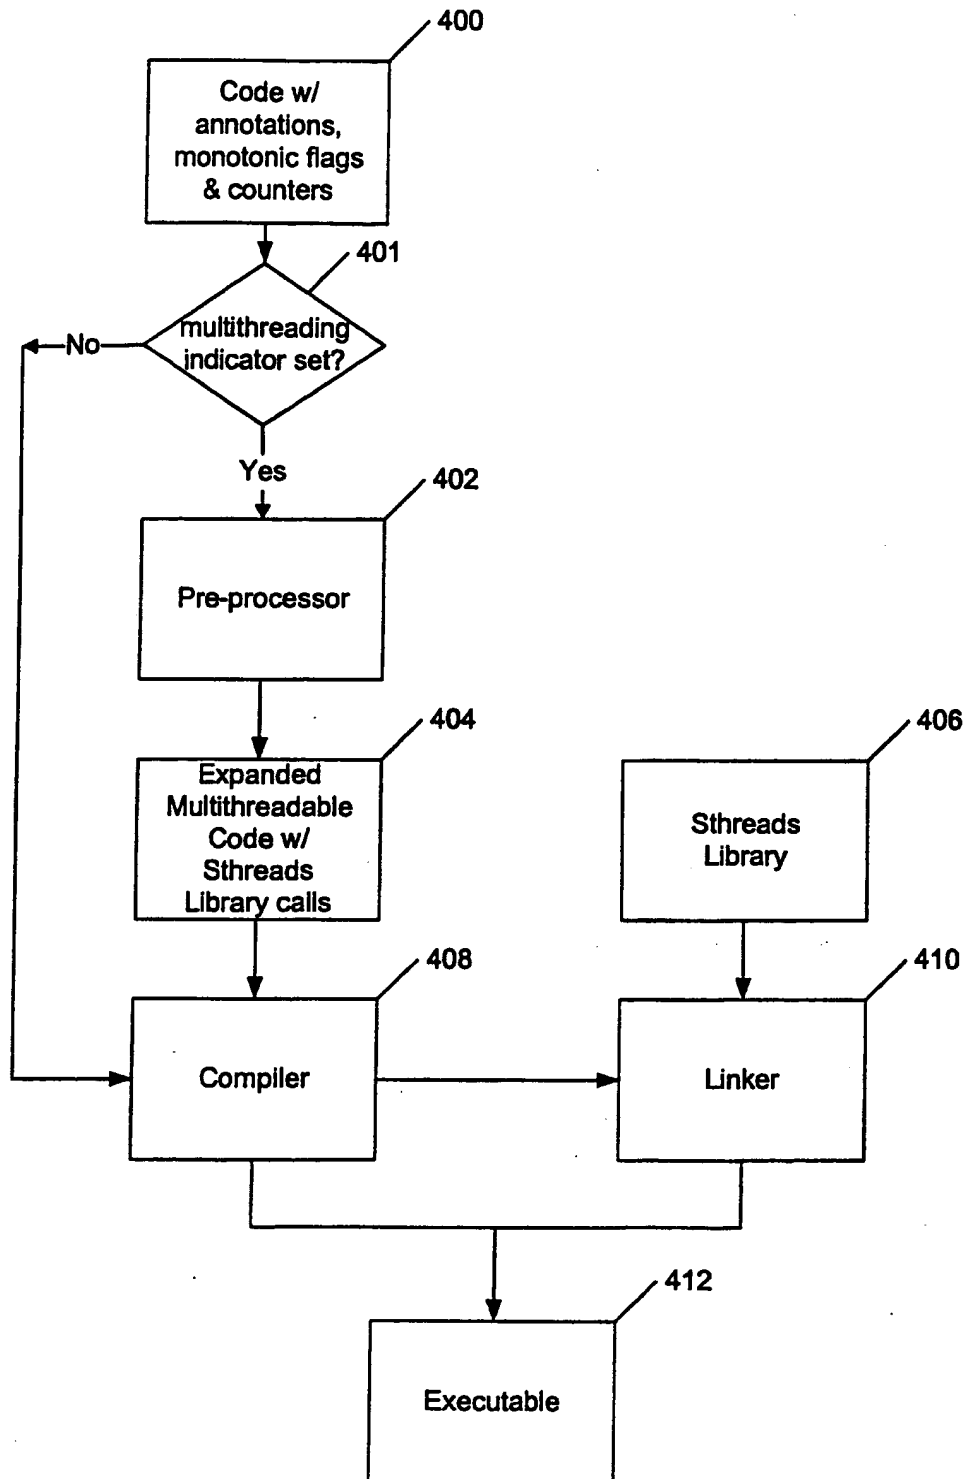


FIG. 24